

Supercomputer Usage Worldwide

**Scientific and Commercial
Benefit is Certain but
Government is Still the
Predominant Sponsor**

LOS ALAMOS NATIONAL LABORATORY



3 9338 00285 3793

Los Alamos
NATIONAL LABORATORY

Supercomputer Usage Worldwide

Harvey J. Wasserman

Margaret L. Simmons

Computer Research and Applications Group

Ann H. Hayes

Advanced Computing Laboratory

LOS ALAMOS NATL. LAB. LIBS.



3 9338 00285 3793

Los Alamos

Los Alamos National Laboratory
Los Alamos New Mexico 87545

LAUR-95-375

February, 1995

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither The Regents of the University of California, the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by The Regents of the University of California, the United States Government, or any agency thereof. The views and opinions of the authors expressed herein do not necessarily state or reflect those of The Regents of the University of California, the United States Government, or any agency thereof.

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7406-ENG-36.

Contents

Executive Summary	1
1. Introduction.....	3
2. Supercomputing Usage Worldwide by the Numbers	7
2.1 Introduction	7
2.2 Installations of Supercomputers Worldwide	10
2.2.1 Effect of Japan's Numerical Wind Tunnel Supercomputer	13
2.3 Supercomputing Usage by Vendor	14
2.4 Supercomputing Usage by Business Sector	31
2.4.1 Categorizing Supercomputer Users	31
2.4.2 Supercomputing Usage by Sector in the United States	35
2.4.2.1 U.S. Government Usage	35
2.4.2.2 U.S. University Usage	43
2.4.2.3 Commercial Supercomputing Usage in the U.S.	44
2.4.2.3.1 In-House Usage in the U.S.....	44
2.4.2.3.2 Supercomputing in Other U.S. Industries	45
2.4.3 Supercomputing Usage by Sector in Japan	50
2.4.3.1 Japanese Government Usage	59
2.4.3.1.1 Government Aerospace Research in Japan	61
2.4.3.2 University Supercomputing in Japan	62
2.4.3.3 Commercial Supercomputing Usage in Japan	63
2.4.3.3.1 In-House Supercomputing in Japan.....	63
2.4.3.3.2 Supercomputing in Other Japanese Industries.....	64
2.4.3.4 Effect of the 1994 Japanese Government Procurements	75
3. Presentation of Case Studies	79
3.1 Introduction	79
3.2 Matsushita Electric Industrial Company.....	81
3.3 Computer Technology Integrator Co., Ltd. (CTI)	85
3.4 Toyota Central Research and Development Laboratory, Inc.	87
3.5 Toyota Motor Corporation	90
3.6 Nissan Motor Corporation	92
3.7 CRC Research Institute, Inc.....	94
3.8 Taisei Corporation	96
3.9 Hitachi, Ltd. General Purpose Computer Division	99
3.10 Angstrom Technology Partnership's Joint Research Center for Atom Technology and Agency of Industrial Science's Research Information Processing Center.....	103
3.11 Grumman Data Systems Corporation	105
3.12 National Center for Supercomputer Applications	106
3.13 Bechtel Corporation.....	108
3.14 Los Alamos National Laboratory Engineering Sciences and Applications Division	109
3.15 Los Alamos National Laboratory Advanced Computing Laboratory.....	111
3.16 Some Conclusions from the Case Studies	113
4. Summary and Conclusions	115
References.....	123
Glossary	129

Contents

1. Introduction	1
2. Theoretical Framework	2
3. Methodology	3
4. Results	4
5. Discussion	5
6. Conclusion	6
7. References	7
8. Appendix	8
9. Glossary	9
10. Index	10

Executive Summary

This report provides a comparative study of advanced computing usage in Japan and the United States as of Spring 1994. It is based on the findings of a group of U.S. scientists whose careers have centered on programming, evaluating, and designing high-performance computers for over ten years. The report is a follow-on to an assessment of supercomputing technology in Europe and Japan that was published in 1993. Whereas the previous study focused on supercomputer manufacturing capabilities, the primary focus of the current work was to compare where, and how supercomputers are used.

Research for this report was conducted through both literature studies and field research in Japan.

The key judgments from the project are summarized below.

- Japanese researchers are as adept at applying supercomputers to science and engineering problems as are researchers in the United States. However, many Japanese supercomputer installations still use old, proprietary, mainframe-based operating environments which make supercomputer research less productive.
- Supercomputer vendors in the U.S. remain largely dependent on U.S. Government-funded purchases, and sufficient expansion into commercial sectors has not yet occurred. In an abrupt change from several years ago, the Japanese government is now the leading supporter of high-performance computing in Japan. Although a recovery from the Japanese recession may reverse this trend, other factors suggest that in neither country can the manufacture of highly advanced supercomputers survive without significant government support.
- Throughout the world, there are many instances in which the use of supercomputing provides significant commercial benefit, and there are some areas in which computing needs will increase. However, because supercomputers require large purchase and maintenance costs, increasingly more companies are electing to use low-end supercomputers or scientific workstations instead. Although new applications of supercomputers are being discovered, the rate of growth is not matching the migration to less expensive machines. We conclude that there will always be many applications for which supercomputers are technically justified, but that economic considerations will dominate, and the demand for high-end supercomputers will continue to shrink.
- Japanese companies currently lack access to advanced parallel computing systems and they lack access to expertise in using those systems. The access afforded companies to state-of-the-art facilities at universities and National Laboratories in the United States gives those companies a distinct advantage over their Japanese counterparts.
- Japan is currently behind the United States in the application of parallel systems to large-scale simulations. Japan's disadvantage is in both applications and system software. However, Japanese

supercomputer vendors, having learned from their experience with vector computing, are now stressing applications development, and significant improvements are expected in the next few years.

- The Japanese government's efforts to stimulate the economy and follow the U.S. lead in establishing formal, large-scale computing programs has resulted in the placement of some very expensive, state-of-the-art, high-performance computing systems at key facilities throughout the country. Although not explicitly declared as such, Japan has a few areas that might be viewed as "grand challenge" scientific computing applications. Foremost among these is aircraft design, which has recently taken a "quantum-leap" in capabilities. Others are nuclear power-related simulations, single-atom and single-molecule behavior, and solid-state physics. Grand challenge computing in the U.S. places more emphasis on development of computing infrastructure than it does in Japan.

1. Introduction

"Supercomputing" may be defined as the use of the most powerful computing systems in existence to carry out calculations to exceptional levels of accuracy on problems requiring immense numbers of operations and/or data. Supercomputers have been used for such calculations for about 20 years and today they are at the leading edge of the information technology revolution that is itself possibly the most important change since the industrial revolution.

One of the most important aspects of supercomputing is the wide range of applications that fit within the generic definition given above. In fact, because the technique of simulating complex phenomena by computers can be applied to nearly every existing field of inquiry, simulation is now regarded as an entirely new field of scientific discovery, separate and distinct from the traditional methods of experimentation and development of theory.

With the ability to model such a wide range of both natural and man-made phenomena, supercomputers can provide more than just a means by which to increase the scientific knowledge base; they also have the potential to benefit virtually all sectors of the economy. Indeed, today, in addition to being commonplace within universities and both civilian and military government agencies, supercomputers have also been installed in many leading commercial entities. The advantage that supercomputing simulation can offer industries is reduced time in design, reduced prototype cost, and improved performance of the final product. A brief list of commercial applications follows [1].

- Automotive: crash worthiness; simulation of airflow in and around a vehicle; combustion analysis in engines; engineering studies of materials;
- Petroleum: analysis of seismic data as part of the exploration process; simulation of oil reservoirs;
- Electronics: simulation of semiconductors and collections of integrated circuits;
- Aerospace: analysis of fluid flow in and around an aircraft body; combustion analysis in engines;
- Manufacturing: analysis of fluid flow in, and mechanical properties of, pumps;
- Chemistry and Pharmaceutical: simulation of refineries and chemical manufacturing processes; simulation of the interaction of potential new therapeutic agents with biological subjects;

The critical role that supercomputers can play fueling innovation in these areas suggests that they are a kind of vital national resource. As such, the well-being of the entire industry that produces supercomputers is in some way related to the well-being of the nation, either in an economic sense or in terms of a more traditional view of national security.

Last year we performed an assessment of supercomputing technology in the United States, Europe, and Japan [2]. Motivated by the relationship between the supercomputer industry and economic security, the focus of the report was the supercomputer *producers*, and the goal was to determine the challenges faced by U.S. supercomputer manufacturers from those in other countries. To answer this question, a wide range of contributing strengths and weaknesses was examined. To a large extent, though, the question boiled down to supercomputer performance, and to which countries could or could not produce the kind of high-performance machines required to carry out the types of simulations described above. Since supercomputer performance is such a well-known and persistently-studied characteristic, and, more importantly, since it is the most important indicator of technological competence, this was a reasonable approach to take.

However, in assessing a nation's overall supercomputing capability, the ability to produce and sell the fastest supercomputers is only part of the equation. An equally, if not more important question is, given that state-of-the-art, globally-competitive supercomputers are available, how are such machines put to use; or more precisely, are there differences in the way supercomputers are used in different countries?

Thus, in this report, we will focus on supercomputer *users*. Of primary concern is whether commercial entities in Japan are more advanced in their application of supercomputers to industrial problems than in the United States.

To answer this question we will explore the distribution of supercomputers in the U.S. and Japan among various business sectors. We will also present case studies of usage in selected areas. Note, however, that there are many computing applications of interest to industry that are not currently regarded as being traditional supercomputing areas, and as such, will be outside the scope of this paper. These would include transportation systems, including air traffic control, computer-aided manufacturing systems, sensor systems for satellites, seismic data, or pollution distributions, patient monitoring systems and other advanced medical equipment such as magnetic resonance imaging, and surveillance or other "real-time," event-driven applications such as transaction processing [3]. We mention these so that by providing this contrast, we can more completely define what we mean by "supercomputing," although we note that it is likely in the future, as computer performance continues to improve, there will be increasing ambiguity between what is and what is not supercomputing.

There are several reasons why there might be differences in U.S. and Japanese supercomputing research. Japan is the only other country in the world that has a well-established indigenous supercomputer industry. In our previous report [2], which included a brief history of the supercomputing industry, we highlighted several notable differences in the way that supercomputing developed in the two countries. Two critical points were: (1) The supercomputer business in the U.S. began largely in support of military and defense programs, whereas in Japan, defense applications are virtually non-existent; (2) Japanese supercomputer manufacturers had already established themselves as dominant suppliers of mainframe and other computing equipment in Japan, and so the introduction of supercomputing frequently came in the form of an upgrade to existing systems. This, we postulate, might have made supercomputing more

attractive to commercial users. In contrast, the dominant U.S. supercomputer supplier had no pre-existing customer base. Therefore, supercomputing in the U.S. might have been viewed as a more exotic technology, one that required a steeper learning curve and possibly longer start-up time until results could be obtained.

These two factors contributed to what was believed to be a more widespread usage of supercomputing by industries in Japan than in the United States. One report several years ago suggested that close to 70 percent of Japanese machines were in the commercial sector, far exceeding the demand from government, research organizations and universities [4]. In the U.S., the market composition may have been nearly the opposite.

attached to the ...
...
...

The ...
...
...

2. Supercomputing Usage Worldwide by the Numbers

2.1 Introduction

We begin our assessment of supercomputer usage worldwide by examining, in terms of a few simple metrics, the distribution of supercomputers throughout the world. The variables in this analysis are the number of installations, computational power, vendor, type of machine, and geographical location. The goals of the analysis are to identify gross trends in usage, such as increased emphasis in one commercial sector in one country relative to another. However, we stress from the outset that although this kind of analysis may yield some interesting trends, it provides merely a suggestion of supercomputing activity and not necessarily a proven, genuine superiority in any one area. There are a variety of reasons for this, which we now explain.

The primary source of information for this phase of the study is the TOP500 list [1], which catalogues data for what is believed to be the 500 most powerful computing sites in the world. We have used the TOP500 list published in November, 1993, and have expanded it to include the new installations of which we are aware, so that the database now comprises a total of 580 sites. The most important new additions are from installations of the new Cray Research Inc. (CRI) massively-parallel system (the T3D) and the installations arising from the recent procurements of the Japanese Government during late 1993 and early 1994. The effect of this latter addition will be discussed separately below.

Before proceeding to the results it is necessary to outline the sources of error associated with such an approach, which are many and varied.

The first level of inaccuracy comes from using the TOP500 list, not necessarily because the data are in error, but simply because it includes only the 500 most powerful systems. In so doing it eliminates many low-end supercomputer systems on which useful computations may be done, such as the CRAY X-MP, some single-processor CRAY Y-MP systems, the newer CRAY Y-MP/EL model, older Fujitsu and NEC models, and numerous Convex systems. The omission of these systems may be quite significant. According to data revealed by Cray Research's Chief Operating Officer Robert Ewald, CRI's most significant growth recently has been in the markets served by its low-end "EL" systems.

In fact, the TOP500 list possibly has an inherent bias toward non-commercial sites because in the U.S., Japan, and Europe, the most powerful computing systems tend to be at national laboratories and universities. The list we used for our research is not by any means intended to be a fully-comprehensive marketing survey of high-performance computing systems. For example, over 1,200 Convex C-2 and C-3 systems have been installed worldwide. However, fewer than 50 Convex systems are included in our database, because the majority of the Convex systems have performance below the minimum cutoff for the TOP500 list.

On the other hand, for the purpose of this report there is good reason to focus on the most powerful of installed systems. This is because we assume an interest only in the most computationally-intensive scientific and engineering kinds of applications which require the most computationally-capable machines.

It is traditionally very hard to obtain data about where supercomputers are being used. The most comprehensive data for Japanese machines comes from lists published by Dr. David Kahaner, of the Office of Naval Research, Asia, in September 1990 [2], February 1992 [3], and December, 1992 [4]. Some of Kahaner's data were obtained from a variety of informal sources such as newspapers. Although the data were reliable when published, as of this writing we have little information about changes that have occurred since then. Because many of the machines in the Kahaner report were low-end systems, it is possible that many sites have upgraded to higher-performance systems since the list was published. However, only a very limited amount of documented update information is available.

Computer vendors are generally reluctant to make available lists of their customers. Occasionally, a customer will prefer that its identity not be publicly known. Furthermore, even with Cray Research, we have had difficulty in obtaining lists of their "public" customers, both in the U.S. and in Japan. We have at our disposal a list of sites belonging to the Cray User Group (CUG) Society, which we have also used to supplement the basic TOP500 list. However, the CUG list may be inaccurate, imprecise, or simply incomplete because some CRI customers choose not to become CUG members. On the other hand, some vendors have been willing to supply data. We obtained a list of NEC supercomputer sites (however, only as of December, 1992) from an NEC employee in Switzerland, and obtained a list of Fujitsu/Siemens customers valid through February 1994 from a Siemens salesperson via electronic mail. (In Europe, Siemens/Nixdorf markets Fujitsu-manufactured supercomputers under its own name. For the purpose of this study we included such machines under the Fujitsu name.)

Other known exclusions from our database include many Hitachi S3600 systems, with peak computational rate between 0.2 and 2.0 GFLOPS, installed during late 1993. It is important to note that our study thus provides a "snapshot" of supercomputing installations as of mid-1994. The user base is constantly changing, and indeed, newer versions of the TOP500 list have been published. However, for the purposes of this report, which attempts to discern *trends* in U.S. and Japanese supercomputing, the older data are probably sufficient. The conclusions section in Chapter 5 contains some general comments on how supercomputing seems to be evolving.

In analyzing trends in the supercomputer market and user base it is tempting to use as a metric the number of active installations. For example, in [5], Johnson and Cavallini concluded that the role of the U.S. Government Laboratories in supercomputing has changed "dramatically" because from 1980 to 1990 their share of *installed systems* dropped from 64% to 28%. However, because virtually all supercomputer systems now available are multiprocessors, estimates based only on the number of installed systems can easily yield misleading results. The difference in peak potential computing power between the least-powerful single processor vector machine and the most-powerful parallel machine is about a factor of

300. Even among single-processor vector systems potential peak performance can vary by about a factor of ten. In the database of installations we shall use for our study we would estimate that fewer than 20% of the machines included are single-processor systems.

Thus, it becomes preferable to also use some actual estimate of an entire system's computing power as a metric. Now the problem becomes how to estimate this power. Although there are significant problems associated with doing so, we will use the LINPACK benchmark results [6] for our performance data. As noted by Dongarra [6], the LINPACK benchmark is not intended to reflect the overall performance of a computing system but rather the performance of a system dedicated to solving one particular kind of computational problem: namely, a dense system of linear equations. For the same multitude of reasons that a supercomputer's theoretical peak performance is never obtained when users run real programs, the performance measured by LINPACK is rarely, if ever, obtained on real programs [7]. It is at best an upper bound on realizable performance, somewhat more accurate than the manufacturer-supplied peak rating; but it should still be regarded as no more than an indication of computing *potential*.

The primary advantages of LINPACK are that it is an experimentally-determined performance measure, as opposed to the paper specifications of a machine and, as noted in [6], it is the only measured performance metric available for all of the machines in which we are interested. Therefore, with adequate caution, we will rely on it as a means of correcting the problems associated with counting installations.

The TOP500 list includes a special version of the LINPACK benchmark that measures the best-possible computing speed on a given machine, by allowing arbitrarily large problem sizes and special programming tricks. This method is used to enable the largest parallel systems to demonstrate their capabilities adequately; however, it again leads to greater differences between the benchmark ratings so obtained and the performance a user might observe on a real program. Therefore, in our opinion, the method generally leads to potentially greater error between LINPACK performance and "real-world" performance for the massively-parallel systems in the database than it does for the vector systems. Partially for this reason, we will generally separate vector and parallel systems when attempting to identify trends in usage.

In summary, counting both the number of supercomputer installations and some measure of supercomputer power are important in characterizing supercomputing usage worldwide. The former gives a rough idea of the supercomputer user base and the availability of supercomputers to different kinds of users. The latter gives a rough idea of how much work can be done at a given site or within a given business sector. To some extent it also suggests how much money an institution is willing to spend on a system.

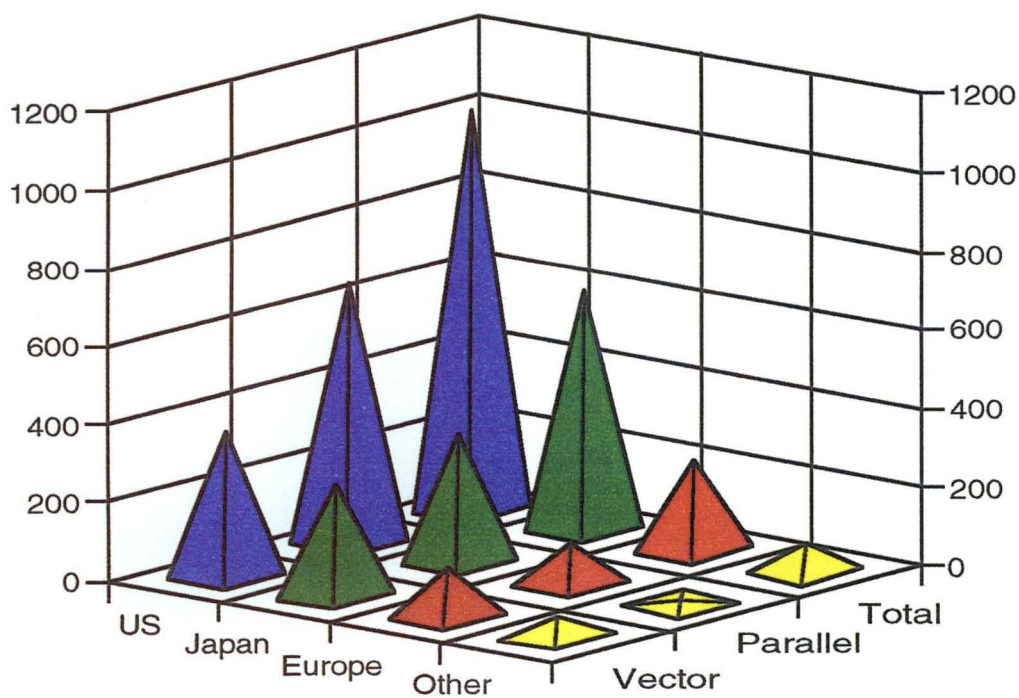
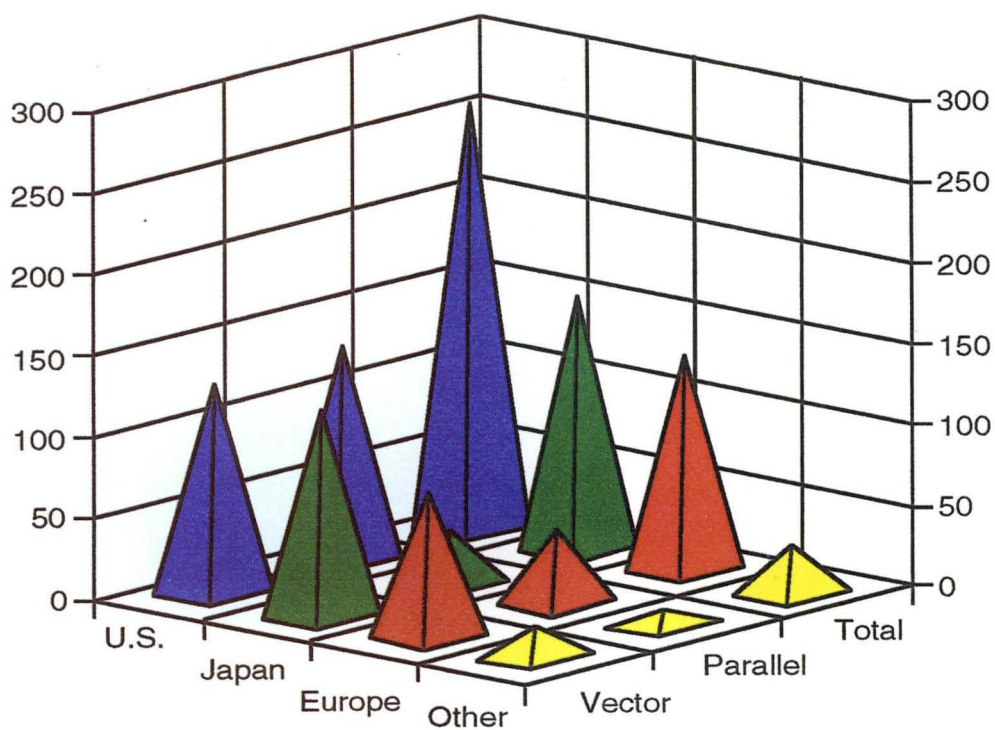
2.2 Installations of Supercomputers Worldwide

In the database of 580 supercomputers the U.S. has about twice as many installations as Japan (Figure 1). The U.S. advantage is due mostly to a much larger number of parallel computing systems. In Japan there are fewer than three dozen parallel systems installed, compared with well over 100 in the United States. The U.S. and Japan have about the same number of vector supercomputers.

Based on LINPACK, supercomputing power in the U.S. is about 60% greater than that of Japan (Figure 2). Again, most of this advantage is from the larger number of parallel machines; the U.S. has about two times as much parallel computing power as Japan. In terms of vector supercomputing power, the U.S. has only about a 25% advantage over Japan. Although we will repeat this message several times during this report, it is important to distinguish between these two kinds of supercomputing power, because we feel that the majority of production computing worldwide is currently carried out on vector supercomputers, whereas most parallel systems are still in research computing environments. "Production" computing means using computer programs that have been "debugged" and scientifically validated and that do not change much on a routine basis, for the purpose of answering a specific science or engineering question. Research computing means development and testing of new techniques largely for the purpose of expanding knowledge about those techniques.

In our FY93 report [8], we stated that "the total high-performance computing power installed in Europe is estimated to be comparable to that of Japan." We know now from our current compilation of data that this was not strictly true during 1993 and it is even further from the truth today. Although in 1993 Europe and Japan had about the same number of vector supercomputer installations, vector computing power in Japan was about twice that of Europe. Today there is much more parallel computing power in Japan than there was when we wrote our earlier report, so that Japan now has over three times as much computing power (LINPACK) than Europe. Japan still leads Europe by a factor of two in vector computing power and it leads by more than a factor of three in parallel computing power. This latter fact is in spite of the number of parallel systems being greater in Europe (45 vs. 31).

The countries in the "Other" category include Australia (11 installations), Brazil (2), Canada (3), Mexico (2), Saudi Arabia (2), Singapore (1), South Korea (4), Taiwan (2), and Peoples Republic of China (1). Of the 28 installations in this category, there are 13 universities, 5 weather prediction facilities, 4 non-defense government installations, 2 sites related to the petroleum industry, one site in the electronics business and one site in the automobile business. One site (in South Korea) is believed to be for defense related work. Incidentally, the NEC SX-3 system installed in Singapore is actually a "pre-owned" system, used between November, 1990 and February 1994 at the Regional Computing Center of the University of Cologne [9]. The university decided to replace the SX-3/11 system with a workstation-based computing system and not to renew the lease for the supercomputer.



2.2.1 Effect of Japan's Numerical Wind Tunnel Supercomputer

In the introduction to this chapter we cautioned that large parallel systems could have an unreasonably large bias on the computing power estimates. This is particularly true in Japan, because of the installation of a single, very powerful supercomputer called the Numerical Wind Tunnel (NWT). The NWT is installed at a Japanese Government institution called the National Aerospace Laboratory (NAL), which is somewhat analogous to NASA in the United States but on a much smaller scale. The NWT was built by Fujitsu, Ltd. in collaboration with NAL, although Fujitsu now markets similar versions of the machine on its own under a different product name [10]. There are 140 processors in the NWT and in early tests it achieved a LINPACK speed of 124 GFLOPS, which is the most powerful rating in our database by about a factor of two over the next highest entry. More recent tests of the NWT show that its LINPACK speed has improved to about 170 GFLOPS. The NWT accounts for 20% of the total supercomputing power in Japan and nearly 40% of the total Japanese parallel computing power.

There are reports [11] that suggest that the NWT is providing Japanese researchers with a resource unmatched in the world and that useful research is being accomplished that cannot be carried out anywhere else. We believe that if it is reasonable to make general statements regarding overall computing power on a national basis, as we are doing in this section, then the NWT must be included in the list of Japanese machines, although with the caveat that a significant fraction of Japan's total computing capability comes from a single machine.

2.3 Supercomputing Usage by Vendor

Certainly the most basic characteristic of a supercomputer beyond its classification as either vector or parallel is its manufacturer. In last year's report [8] we described some of the many differences in design between U.S. and Japanese supercomputers, both vector and parallel.

Figures 3 and 4 show the worldwide distribution of supercomputer installations and LINPACK supercomputer power by vendor. The "Other" category here includes systems by KSR, Convex, Meiko, Parsytec, MasPar, IBM, and NCube. Worldwide installations by American vendors (about 78%) far outnumber those from Japanese vendors, and only seven installations with European-manufactured systems are included in the database. Cray Research, with 247 installations, is still the dominant worldwide supplier; but its percentage of total installations has been reduced recently, and the data allow some estimate of the source of this reduction. Installations due to other American suppliers, mostly parallel systems from Intel Corporation and Thinking Machines Corporation (TMC), total 200 systems, 34% of the total. In comparison, installations due to Japanese suppliers combine to 127 systems, 22% of the total. Thus, one can argue that at least until now, greater competition to Cray's market share has come from its own domestic competitors than from overseas competition.

Vector machines are still the dominant type of computational tool, with 363 installations, compared with 218 installations of parallel systems. Figures 5 and 6 show the worldwide distribution of vector supercomputers and their corresponding LINPACK power. Cray Research's dominant share of the *vector* computing market has been reduced relative to previous years, and this has come about largely from the Japanese competitors. However, we note two points in this regard. First, there is little reason to believe that Cray's share of vector systems will erode significantly below the two-thirds level it is now, as we suggested in last year's report. Second, although CRI's share probably will not decrease, perhaps the more important question is what will happen to the absolute size of the vector market. This matter is covered in more detail in a later chapter of this report, but generally speaking, there is a good probability that it will experience significant reduction.

The market for parallel systems is shown in Figures 7 and 8. This area is significantly more dynamic than the vector realm, with numerous companies entering as well as departing. Two important newcomers are Cray Research and Fujitsu, both of which have several of their first-generation parallel systems included in the TOP500 database. Because it is still early in the life-cycle of these products, it is difficult to predict how market share will change. In terms of share of installed systems, i.e., data represented in Figure 7, there is reason to believe that Fujitsu will be the more successful of the two. This is because Fujitsu's parallel system can be sold in small, (relatively) inexpensive configurations, thus leading to more

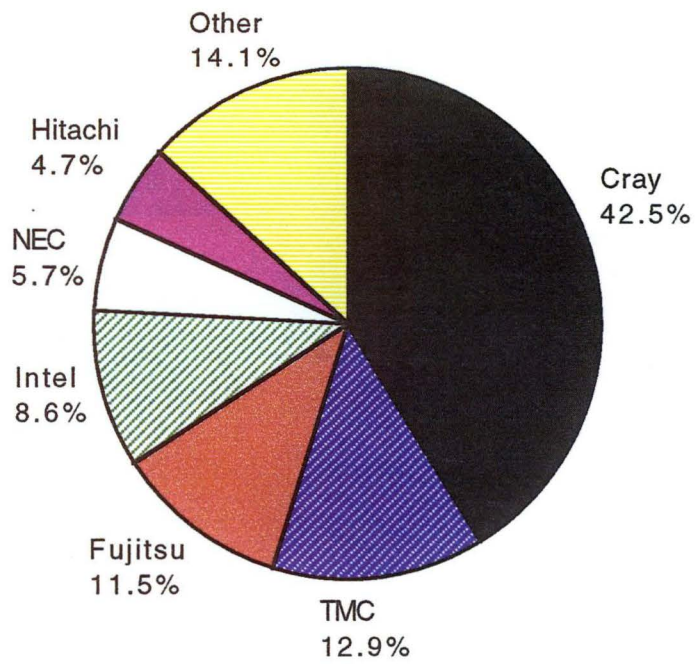


Figure 3. Worldwide Installations of Supercomputers by Vendor (580 systems).

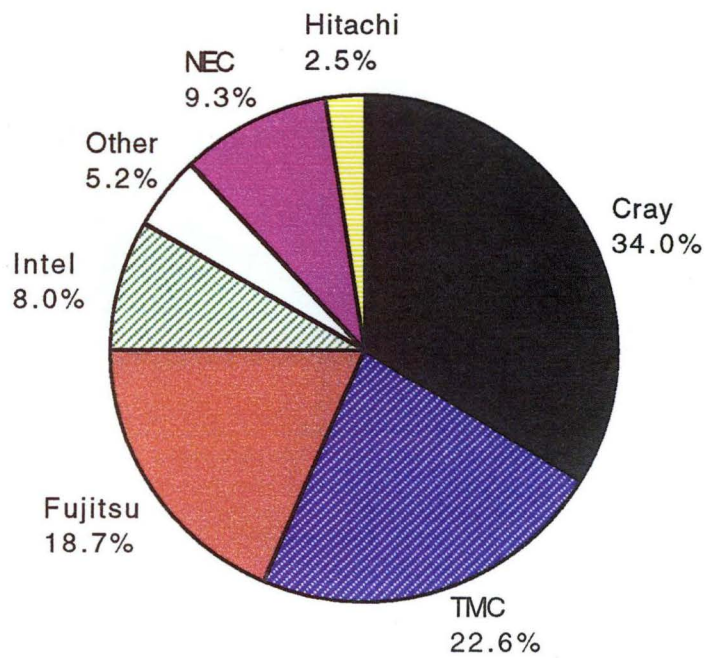


Figure 4. Worldwide Supercomputing Power by Vendor.

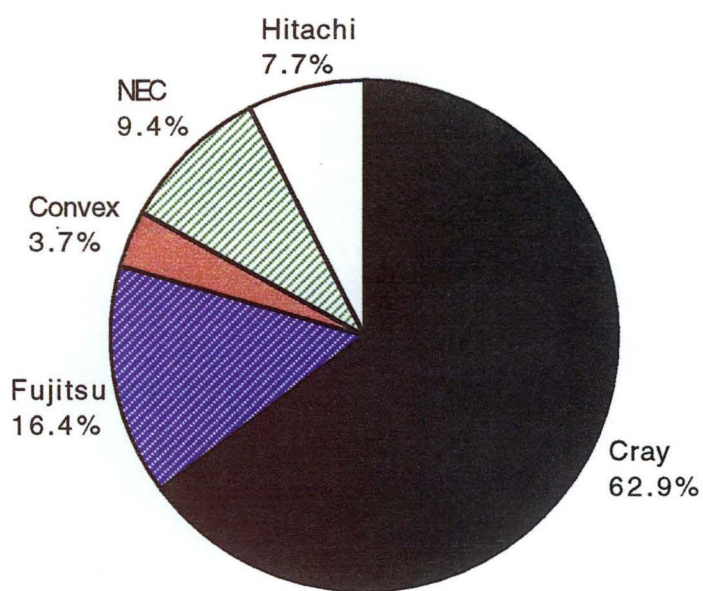


Figure 5. Worldwide Installations of Vector Supercomputers by Vendor (363 systems).

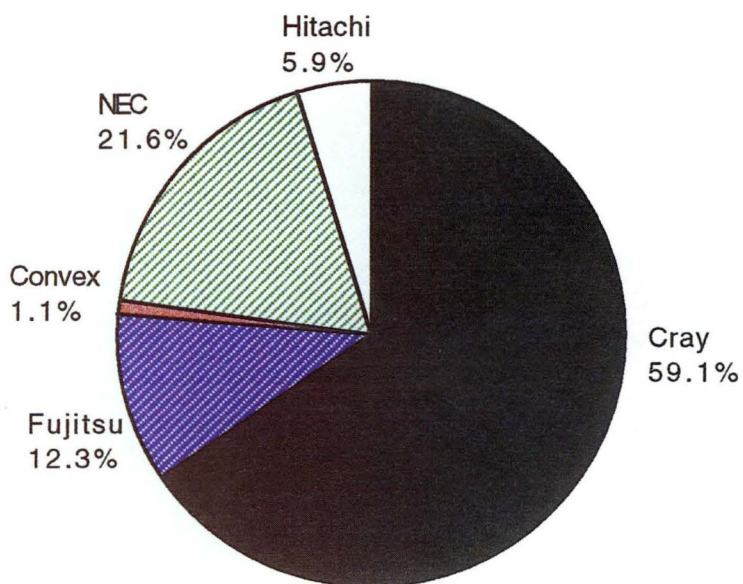


Figure 6. Worldwide Vector Supercomputing Power by Vendor.

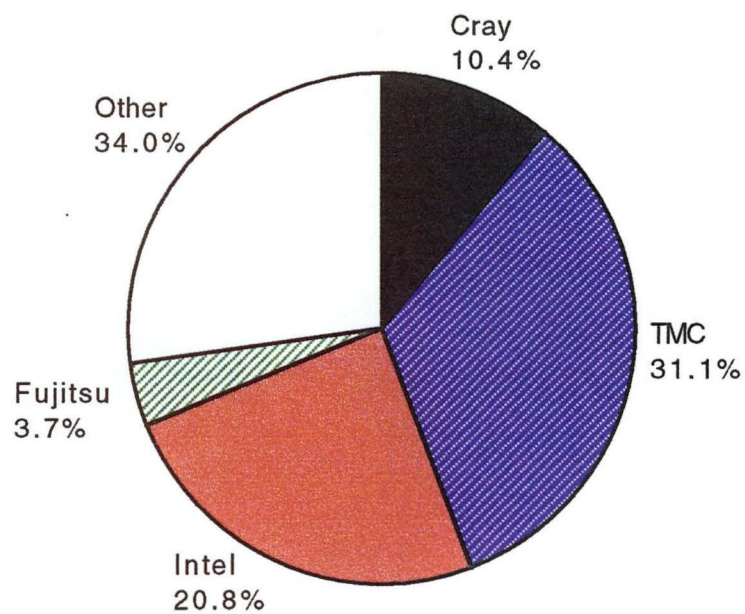


Figure 7. Worldwide Installations of Parallel Supercomputers by Vendor (218 systems).

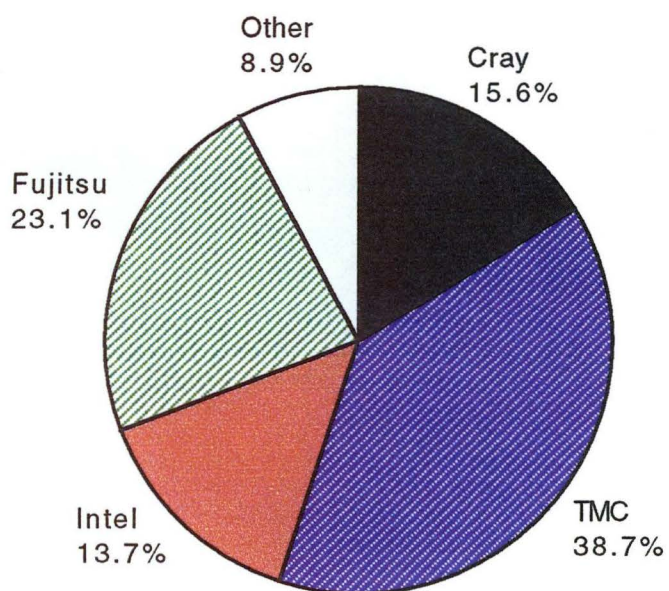


Figure 8. Worldwide Parallel Supercomputing Power by Vendor.

installations. The question of what one can do with these smaller configurations is, of course, a matter that the installations-only data do not comment upon. Besides Cray and Fujitsu, the most important growth in the parallel market will come from IBM. Other Japanese vendors are expected to enter with parallel products soon, but we feel that it is unlikely that these vendors will capture significant market share, at least in the near term.

Comparing Figures 7 and 8, which show installations and LINPACK power of the world's parallel systems, gives the most dramatic example so far of how parallel machines can alter the balance of computational power and why it is therefore misleading to draw conclusions based only on the number of installations. Specifically, Fujitsu accounts for only 4% of worldwide installations of parallel systems with 9 sites, but because of the high potential computing rate of Fujitsu's VPP500 system and the NWT, Fujitsu's share of worldwide parallel computing power is much higher, at about 23%. The computing power of the Intel systems and those in the "Other" category occupy a proportionately smaller share of the total relative to their number of installations.

Figure 9 shows the distribution of supercomputer installations in Japan. The data in this figure show a very different result from the last such compilations of which we are aware. In 1991 and 1992, Furutsuki [12] and Kahaner [4], respectively, presented data on supercomputer installations in Japan obtained from the Japanese journal *Nikkei Computer*. As of the end of 1991, according to those reports, out of a total of 125 supercomputers in Japan, Fujitsu's share was about 50% of the installations, Hitachi was second with about 19%, and Cray and NEC followed with about 16% and 13%, respectively.

The most significant change relative to 1991, then, is that Fujitsu's share is now about 25%, about one-half of what it was in the previous study. However, a direct comparison between the *Nikkei Computer* data and the present data is somewhat misleading. The first reason for this is that the older survey included a great many machines, especially by Fujitsu, that are not included in the current study because they are old and are generally below the minimum performance level for our database. Such machines include many versions of Fujitsu's first-generation supercomputer, the VP-100. These machines may still be in use in Japan (it is not known for certain), but their performance is on a par with today's RISC workstations, which are omitted for practical purposes. It is known that several Siemens versions of these machines (see above) are included in the database, but for these we obtained definite confirmation of their use by Siemens personnel [13].

The second factor contributing to the difference between the current compilation and that of previous reports is the installation of 31 parallel systems in Japan by Thinking Machines, Cray Research, Intel, MasPar, KSR, Fujitsu, and NCube, which now collectively account for about 18% of Japan's total installations. No parallel systems were included in the 1991 reports. The reduction in Fujitsu's share of the installed base since 1992 is counter-balanced by this increase in parallel systems and also by NEC's share of the total installations in Japan, which increased by about five percentage points.

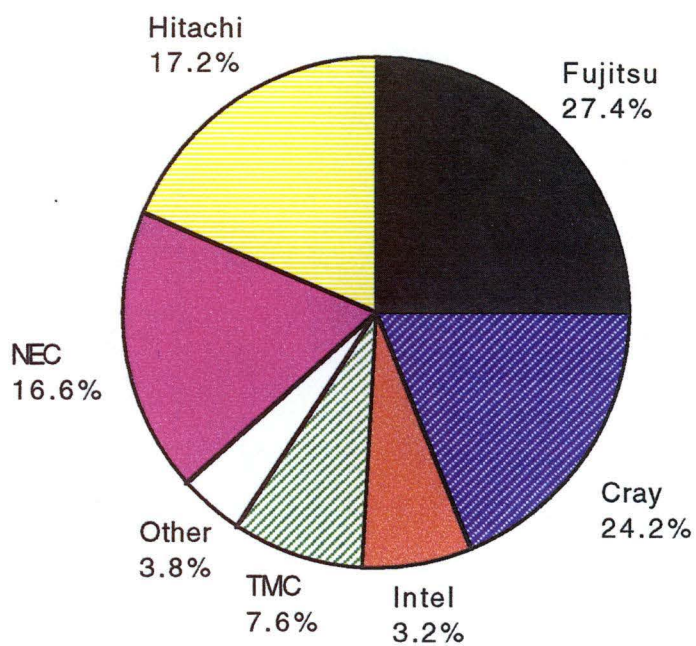


Figure 9. Installations of Supercomputers in Japan by Vendor (157 total).

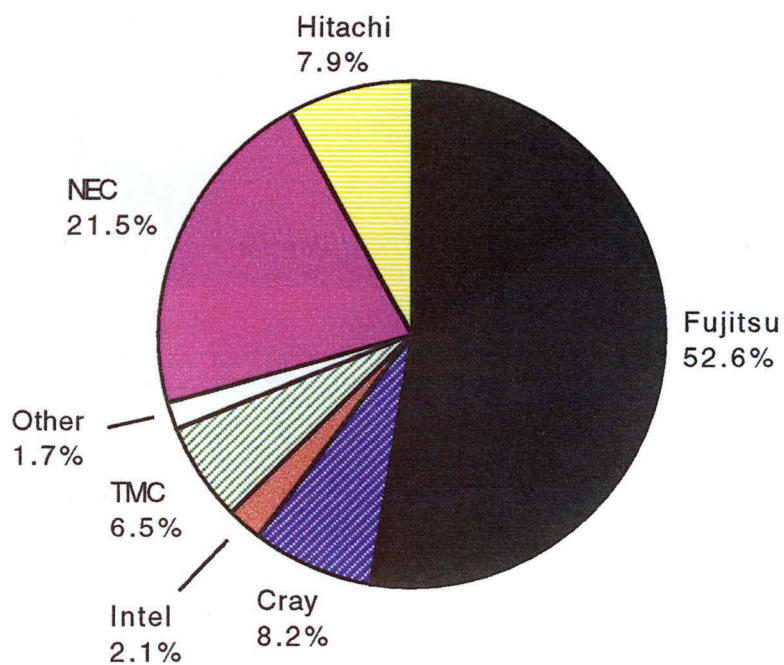


Figure 10. Supercomputing Power in Japan by Vendor.

Because of the omission of lower-end systems in the current study, there is an equal distribution of vector supercomputer installations in Japan among the four vendors (Cray, Fujitsu, Hitachi, and NEC, although note that as mentioned above, our database may under-count Hitachi systems somewhat). Again, this is significantly different from the previous study, in which Fujitsu had a significant lead over Hitachi, NEC, and Cray Research. However, while the installations are equally distributed among the vendors, the vector supercomputing power in Japan is distinctly unbalanced, with NEC claiming nearly one-half of the total. We can see two reasons for this. NEC's advantage over its Japanese competitors is that it is the Japanese manufacturer that has been offering multiprocessor vector supercomputers for the longest time. Fujitsu's vector machines are all single-processor and Hitachi just very recently started selling multiprocessor machines (there are only three in our database). NEC's advantage over Cray Research is that NEC's individual processors are over five times faster than individual processors of Cray's machines (based on LINPACK). Also, the majority of Cray Research vector machines in Japan are not fully-configured systems; all but two have fewer than the maximum number of processors per system. Interestingly, in our database there happen to be exactly the same number of NEC and Cray Research vector supercomputers in Japan (26). However, the average computing power of an NEC vector supercomputer system in Japan is about 5 GFLOPS while the average for a Cray system is about 1.3 GFLOPS.

In terms of combined (vector plus parallel) supercomputing power in Japan (Figure 10), Fujitsu is by far the leader, and its 25% of Japan's installations account for more than 50% of Japan's supercomputing power. Figures 11 and 12 show parallel installations and parallel computing power in Japan, respectively. The latter especially shows that Fujitsu's overall lead comes almost entirely from its overwhelming lead in parallel supercomputing power. Nearly 80% of the parallel computing power in Japan is provided by Fujitsu, one-half of the total power being from the NWT.

For comparison purposes, we show in Figures 13 and 14, installations of supercomputers and supercomputing power in Europe. Cray Research has the dominant share of both, with 59 out of a total of 133 systems.

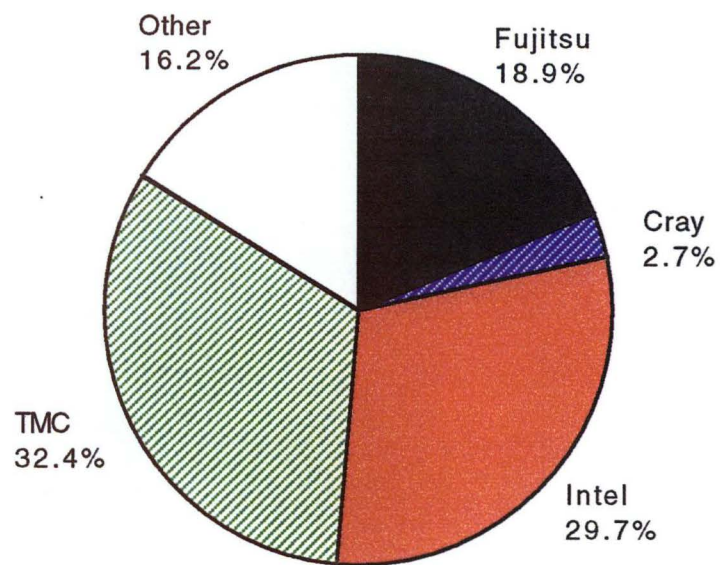


Figure 11. Installations of Parallel Supercomputers in Japan by Vendor (31 systems).

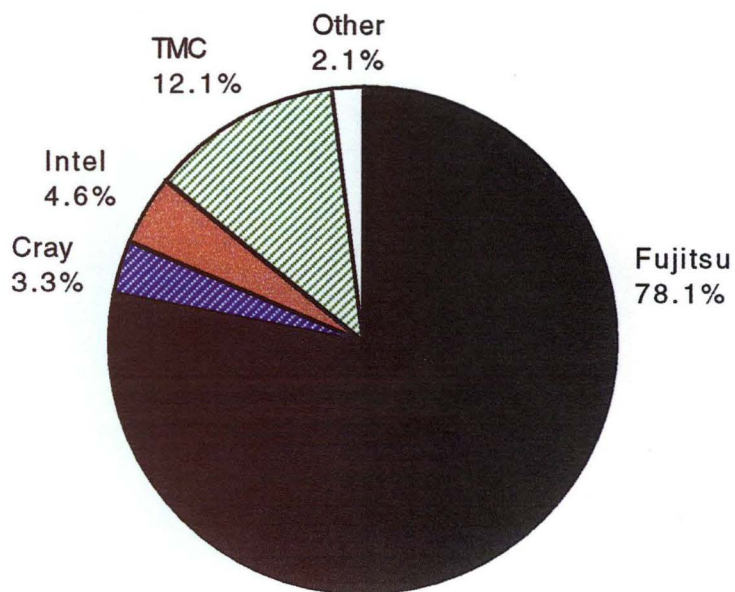


Figure 12. Parallel Supercomputing Power in Japan by Vendor .

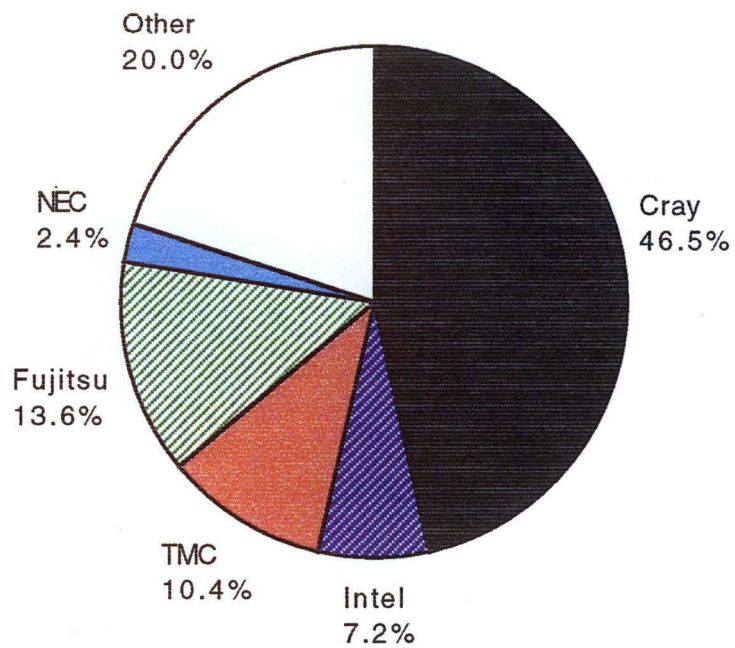


Figure 13 Installations of Supercomputers in Europe by Vendor (133 systems).

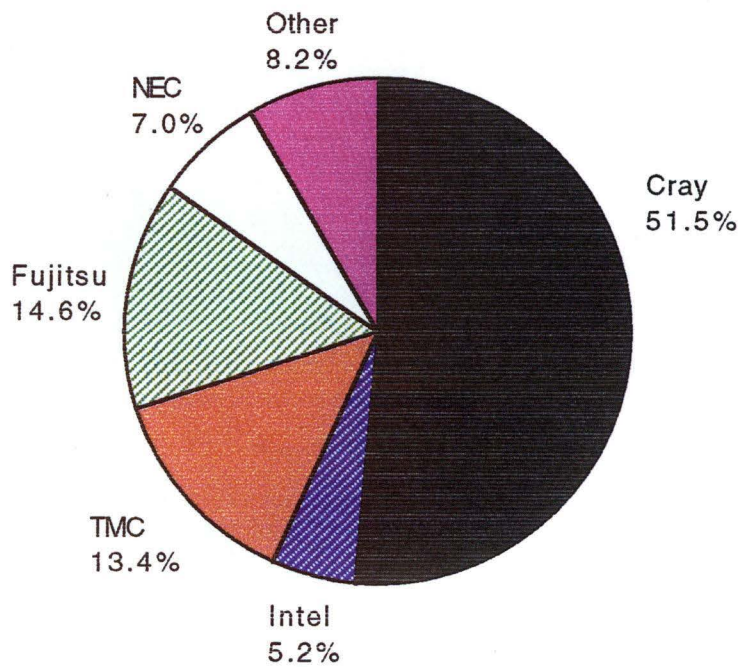


Figure 14. Supercomputing Power in Europe by Vendor.

2.4 Supercomputing Usage by Business Sector

2.4.1 Categorizing Supercomputer Users

The question of where supercomputers are used is the primary focus of this report. To help answer this question we categorized each site in our database according to our estimate of the field of application. Of course, assignment to one category or another contains a certain level of arbitrariness and we will be careful to point out instances in which the categorization is imprecise for one reason or another.

To a first approximation, three categories of usage are important: government, university, and industry. Each will be explained in detail below, and it is important not to base final conclusions on this simple classification. Figure 15 compares the distribution of supercomputer installations in the United States, Japan, and Europe according to these three basic categories, and Figure 16 shows the comparison based on supercomputing power according to LINPACK. The data so far may be summarized as follows:

- The U.S. has many more supercomputers in use in government than Japan does, but the U.S. and Japan have about the same number of supercomputers installed in industry. U.S. industries have about 60% more supercomputing power than do Japanese industries. However, the key point is that the composition of both the government and industry sectors is very different in the U.S. and Japan, as will be discussed below.
- In the U.S. there are about the same number of industrial and government installations. However, *supercomputer power within the government sector is twice that of the industrial sector.*
- In Japan this same situation holds, but it is even more exaggerated. There are even more supercomputers in industry than in government, but supercomputing power within the government sector is still twice that of the industrial sector.

1. The Government of the United States of America

2. The Government of the United Kingdom of Great Britain and Northern Ireland

3. The Government of the United States of America, the Government of the United Kingdom of Great Britain and Northern Ireland, and the Government of the United States of Canada, have agreed to the following terms of reference:

4. The Government of the United States of America, the Government of the United Kingdom of Great Britain and Northern Ireland, and the Government of the United States of Canada, have agreed to the following terms of reference:

5. The Government of the United States of America, the Government of the United Kingdom of Great Britain and Northern Ireland, and the Government of the United States of Canada, have agreed to the following terms of reference:

6. The Government of the United States of America, the Government of the United Kingdom of Great Britain and Northern Ireland, and the Government of the United States of Canada, have agreed to the following terms of reference:

7. The Government of the United States of America, the Government of the United Kingdom of Great Britain and Northern Ireland, and the Government of the United States of Canada, have agreed to the following terms of reference:

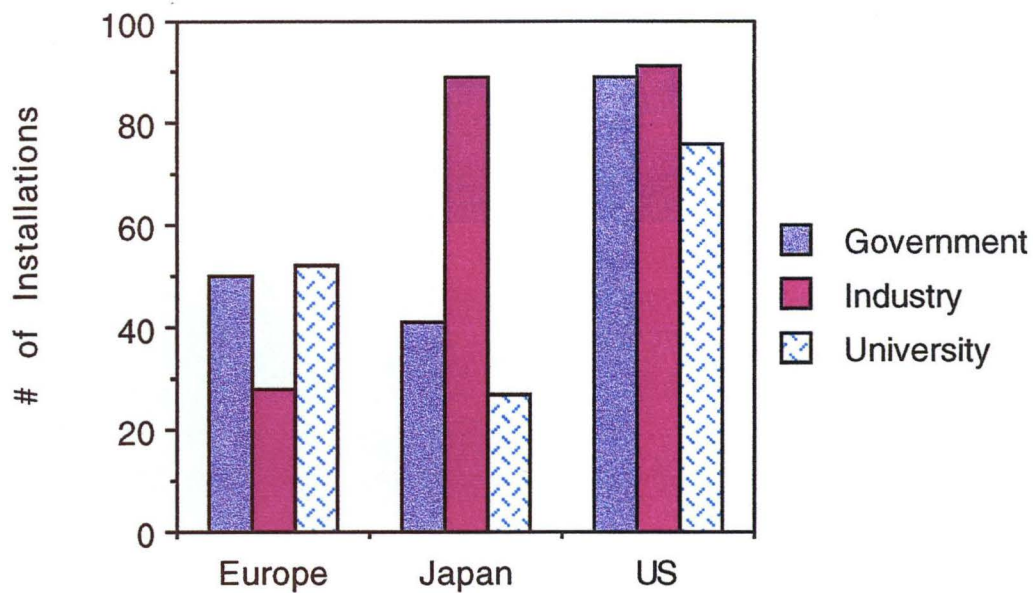


Figure 15. Comparison of Worldwide Supercomputer Installations Based on Three Categories.

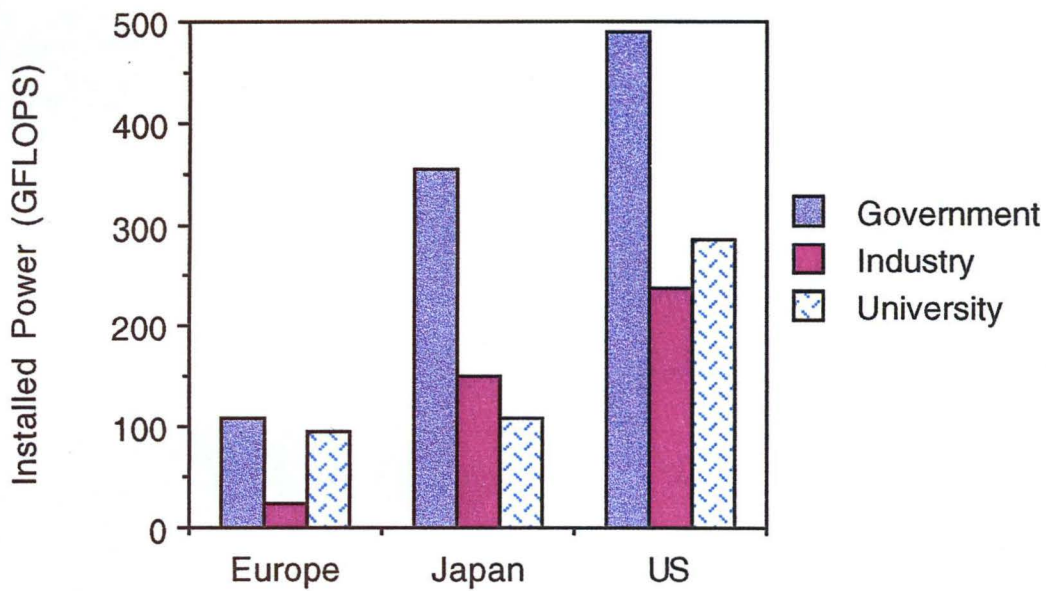


Figure 16. Comparison of Worldwide Supercomputer Power Based on Three Categories. Units are GFLOPS.

2.4.2 Supercomputing Usage by Sector in the United States

Although these results are important, even more revealing differences between the U.S. and Japan are obtained by separating the three simple categories into additional sectors. First, we will discuss U.S. supercomputing usage, sector-by-sector, and then compare with Japan in a later section. Figure 17 shows a detailed view of the distribution of supercomputer installations in the United States from the 580-supercomputer database and Figure 18 shows the distribution of computing power according to LINPACK. Figures 19 and 20 show the sector distribution of vector supercomputer installations and vector supercomputing power in the U.S., and Figures 21 and 22 depict the corresponding sector distribution for parallel supercomputers.

2.4.2.1 U.S. Government Usage

In the U.S., it is particularly important to distinguish between defense related and non-defense related government computing, although this distinction is not easily made in some cases. For example, the Los Alamos National Laboratory (LANL) is assigned the "Government Defense" category because this is the main mission of the Laboratory; however, a significant amount of non-defense related high-performance computing is also carried out at LANL. The same is true for Lawrence Livermore National Laboratory; although in contrast, the Argonne and Oak Ridge Laboratories are assigned to the Government Non-Defense category. Two other separately-defined government categories are aerospace, which in the U.S. consists of 17 supercomputers at seven different NASA sites, and weather prediction, which is carried out at both military and civilian institutions.

When combined, these four government categories account for about one-third of supercomputer installations in the United States. When viewed according to computing power from LINPACK, their share is even larger, accounting for fully one-half of the power. The government sector has 40% of all vector supercomputers installed in the U.S., and 66% of the available vector supercomputing power. Government sites do not have the largest share of parallel supercomputer installations, although they are a close second; and they do have the lion's share of parallel supercomputer power.

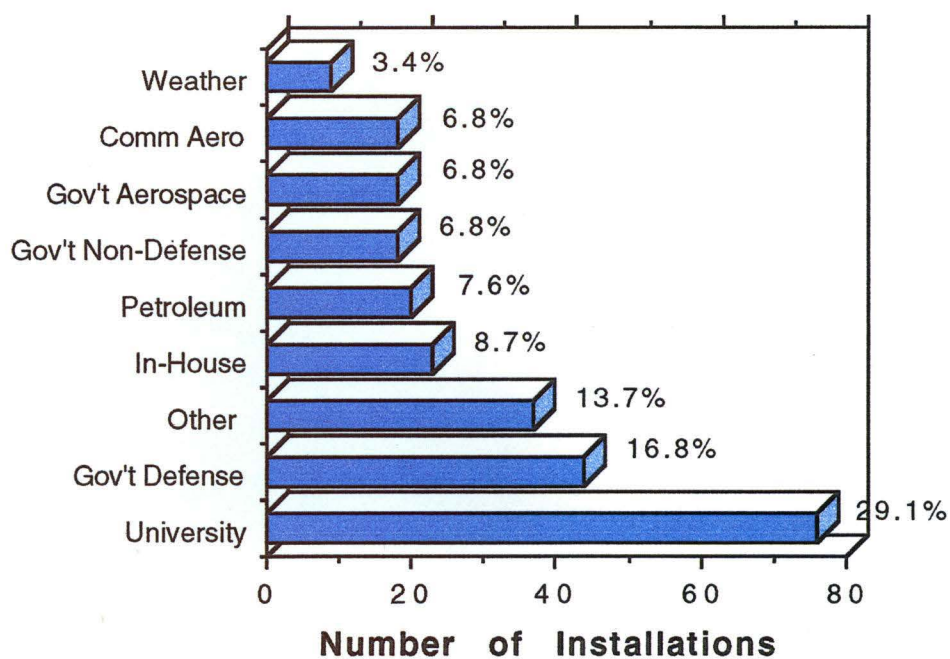


Figure 17. Detailed Sector Distribution of Supercomputer Installations in the U.S. (263 systems).

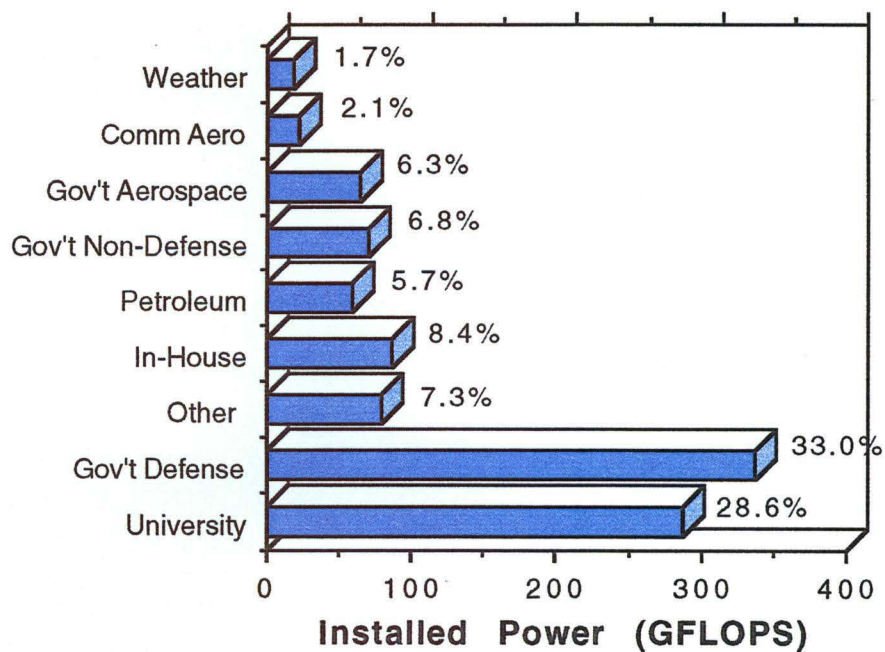


Figure 18. Detailed Sector Distribution of Supercomputing Power in the United States.

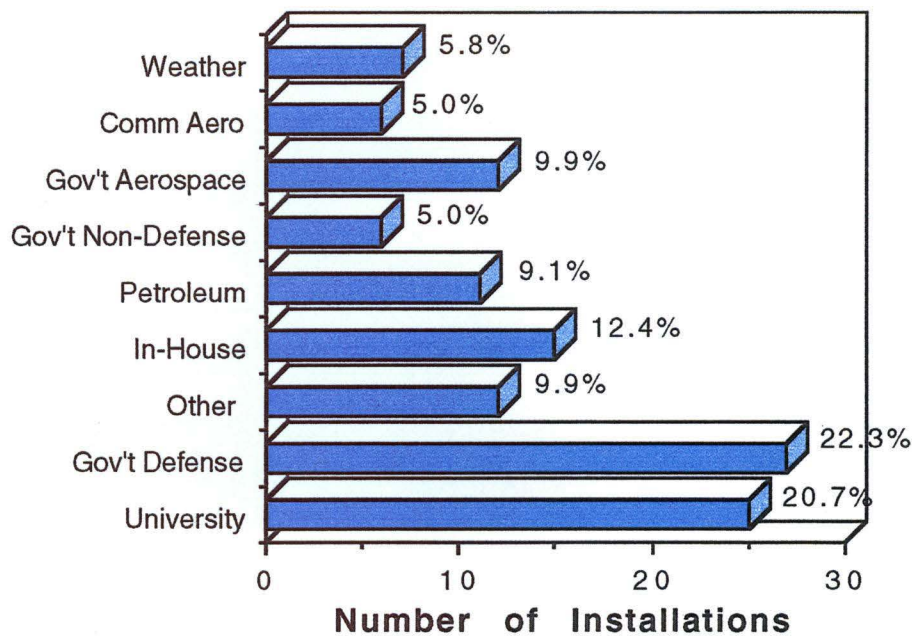


Figure 19. Detailed Sector Distribution of Vector Supercomputer Installations in United States (130 systems).

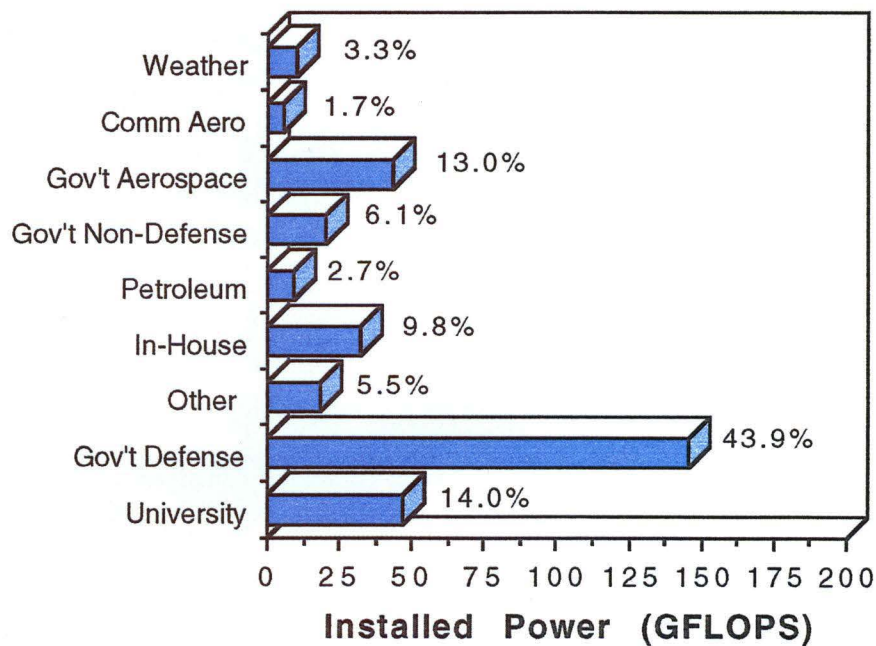


Figure 20. Detailed Sector Distribution of Vector Supercomputer Power in the United States.

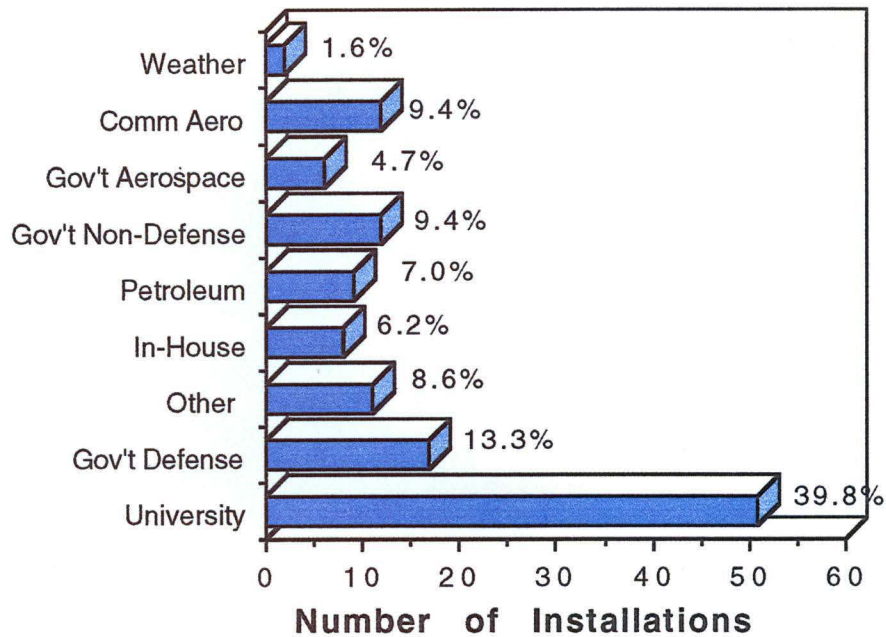


Figure 21. Detailed Sector Distribution of Parallel Supercomputer Installations in United States (133 systems).

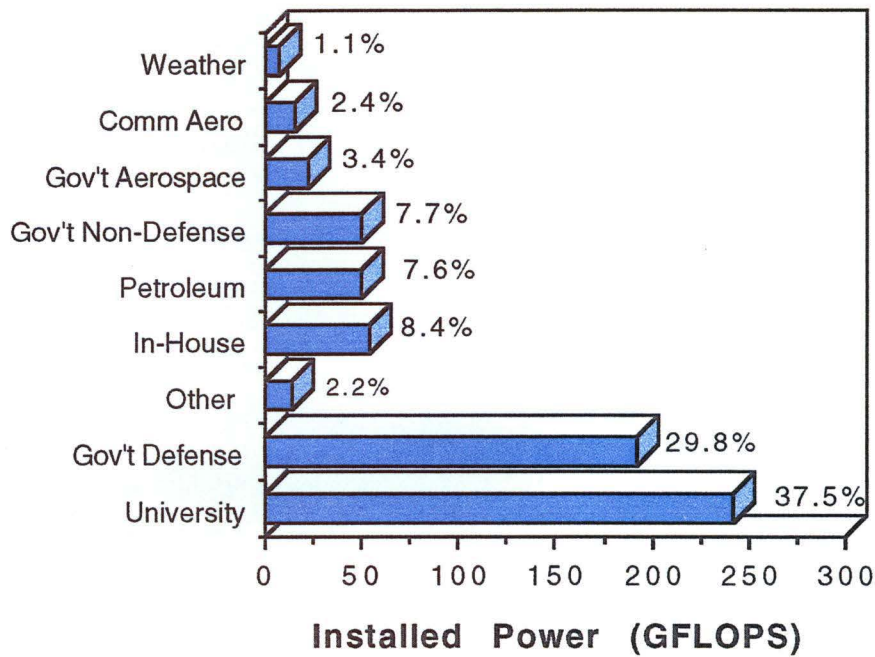


Figure 22. Detailed Sector Distribution of Parallel Supercomputer Power in the United States.

Government non-defense high-performance computer sites in the U.S. are listed in Table 2.1. Whereas users within the defense sector overwhelmingly prefer vector supercomputers, users within the "government non-defense" sector seem to prefer parallel rather than vector machines. Of the 18 machines in the government non-defense list, only 6 are vector machines, three of which are at one site, the National Energy Research Supercomputer Center (NERSC) located in Livermore, California. To a certain extent, this preference for parallel machines represents a new contingent of high-performance computing researchers within the U.S. government. Traditionally, such research was carried out largely at Los Alamos and Livermore, but now Argonne and Oak Ridge National Laboratories are carving out research niches for themselves by investigating the newer parallel machines. These latter two laboratories together have five of the parallel machines in the list. However, the list also suggests a group of both traditional research and non-research institutions that are willing to experiment with newer parallel technologies. The usage of the massively-parallel Connection Machine at the FBI, which is believed to be used as a database engine, is a case in point.

2.4.2.2 U.S. University Usage

After the combined government usage, the second largest category in the U.S. is university supercomputing, which constitutes almost one-third of the total installations (76 machines at 44 different sites) and slightly less than one-third of the total power. Most of the parallel supercomputing systems in the United States are installed at universities (Figure 22). However, it is important to note two points about the U.S. university sector. First, computing strength at American universities was significantly

Table 2.1: Government Non-Defense Supercomputing Sites in the United States

Site	Vendor	Model	Type	LINPACK Rating (GFLOPS)
Argonne National Laboratory	IBM	SP-1	P	3.9
FBI	TMC	CM-5	P	15.1
FermiLab National Laboratory	Intel	iPSC	P	0.6
Jet Propulsion Laboratory	TMC	CM-5	P	1.9
	Cray	Y-MP/1	V	0.3
	CRAY	T3D	P	10.6
Nat'l. Energy Research Supercomputer Center	Cray	CRAY-2	V	1.4
	Cray	CRAY-2	V	2.2
	Cray	C90	V	13.7
National Cancer Institute	Cray	Y-MP	V	2.1
	MasPar	MP2216	P	1.6
National Institutes of Health	Intel	iPSC	P	2.6
Nat'l. Inst. Standards and Tech.	Cray	Y-MP2	V	0.6
Oak Ridge National Laboratory	Intel	iPSC	P	2.6
	Intel	Paragon	P	2.0
	KSR	KSR1/64	P	1.8
	Intel	Paragon	P	15.2
Superconducting Super Collider	Intel	iPSC	P	1.4

augmented during 1994 with the installation of five CRAY T3D systems. The T3D is a massively-parallel system from Cray Research with a high peak potential computing rate. The five T3Ds account for about 20% of all university sector computing power. Second, some of the sites in the database categorized as "University" are also major national centers of computation with a significant amount of accumulated computing power beyond that which the university uses or pays for. The most important example of this is the Minnesota Supercomputer Center Incorporated (MSCI), at the University of Minnesota, which owns seven machines in the database, also amounting to about 20% of the U.S.'s university computing strength. Other important examples of these centers are two National Science Foundation (NSF)-sponsored institutes: Pittsburgh Supercomputing Center, which is jointly run by the University of Pittsburgh, Carnegie Mellon University, and Westinghouse Electric Company, and the National Center for Supercomputing Applications at the University of Illinois. Five NSF centers and MSCI together account for one-third of the university computing sector and, incredibly, 56% of the university sector computing power.

2.4.2.3 Commercial Supercomputing Usage in the U.S.

2.4.2.3.1 In-House Usage in the U.S.

The "In-House" category refers to supercomputers installed within the company that manufactures them. Such machines are used for customer demonstrations, software development both by the manufacturer and by outside software vendors, advanced R&D related to new products and more recently, for selling timesharing services. There is also some basic scientific research that goes on within these companies, both in the U.S. and in Japan. There has arisen a tradition, probably born from IBM, in which a supercomputer company hires world-renowned scientists to demonstrate the power of its machine in solving problems in the scientists' field of research. In so doing the company hopes that some kind of critical scientific breakthrough will be accomplished using their machine, thus producing an important source of publicity.

In the U.S., most important supercomputer vendors, including Intel, Cray Research, KSR, and Thinking Machines, as well as Fujitsu America, Inc., have significant in-house supercomputing capabilities. (Cray Computer Corporation of Colorado Springs Colorado, does not have an entry in our database, although they may have several systems.) In fact, the in-house sector is the largest user of supercomputers in the U.S., after the government and universities. Cray Research, Inc. uses an astoundingly large number of its supercomputers in-house. Our database has 15 Cray systems of various types installed within the company. (The exact configuration and available power of "in-house" supercomputers is subject to error because the systems are often reconfigured to meet particular customer demands.)

Mostly because of the large number of machines installed at Cray Research, the supercomputer manufacturing sector is the largest commercial consumer of supercomputers in the U.S., with about 23 machines total. There is more supercomputing power installed in-house than there is in the petroleum sector, the automobile sector, or the government non-defense sector, to name a few. However, for the

purposes of this report, we suggest that it is probably more important how many supercomputers are being purchased by non-vendor companies with the aim of enhancing the competitiveness of the goods or services that they sell. The supercomputer industry, with about \$2 billion in annual sales [14], is much smaller than that of all the other commercial sectors that use these machines, and so the extent to which a supercomputer vendor uses machines to increase its profitability is probably less significant than the others.

Supercomputing in Other U.S. Industries

Not including in-house machines, supercomputer resources associated with U.S. industries (67 machines) amount to only about 25% of the total installations in the U.S. and only about 15% of the total installed power. The application fields of these U.S. commercial users (including the in-house users) are listed in Table 2.2.

The petroleum industry is the largest non-vendor commercial user in the U.S., with a total of about 20 supercomputers in use providing about 58 GFLOPS total performance. This is nearly twice the computing power installed within the commercial aerospace or automobile industries. However, about one-third of the petroleum industry's available supercomputing power is installed at one company (Exxon), in the form of a 256-processor massively-parallel T3D system from Cray Research. Exxon was one of the first institutions of any kind to install the new Cray parallel system, and the company believes that by using the system for a few months or a year before any of its competitors it can gain a significant advantage over them. So far, only one other petroleum company, Phillips, has followed Exxon by

Table 2.2: Commercial Supercomputing in the United States.

Application	Number of Machines		Combined LINPACK Rating (GFLOPS)
	Vector	Parallel	
Aerospace	12	6	21.3
Automobile	8	0	29.7
Biotechnology	0	1	2.6
Chemical / Pharmaceutical	6	0	7.0
Electronics / Telecommunications	3	4	6.0
Financial	1	3	14.1
In-House	15	8	86.5
Manufacturing	0	1	1.4
Petroleum	11	9	58.6
Service	1	1	10.9
Total	57	33	238.1

ordering a T3D. However, the petroleum industry has been relatively eager to embrace parallel computing in general; nine of the 20 petroleum industry machines in our database are parallel systems, although none come close to the Exxon T3D in terms of potential performance.

The supercomputer computational workload at petroleum companies consists largely of two elements. The first is seismic processing which is used for exploration, and the other is reservoir simulation, which is used to increase the productivity of known oil fields. The large T3D system at Exxon is used primarily for seismic processing and so is the Thinking Machines CM-5 at Mobil. (Interestingly, Mobil used a general-purpose mainframe system for this work until the CM-5 became available about three years ago.) In contrast, reservoir simulation with massively-parallel systems is viewed as a research activity, at least at Mobil.

Seismic computing presents a number of interesting challenges not found in many other supercomputing domains. On one hand, it is highly input/output-dependent – one computer simulation may require hundreds, if not thousands, of computer tapes to be read. The requirement for enormous tape-processing capability plus a virtually insatiable need for more raw computational power to handle all the data suggests that petroleum companies will always be among the most important and consistent supercomputer customers, and indeed, a major part of the driving force toward massively-parallel computing. On the other hand, seismic processing is also by nature a distributed computing discipline. Much of the processing needs to be done at the potential oil fields themselves, rather than at a centralized research laboratory. Sometimes this even means having a computer onboard a ship or on a truck. For this reason, smaller and easier-to-install (air-cooled) systems such as those made by Convex have become popular.

The petroleum industry is effectively the largest user of Japanese-manufactured supercomputers in the United States. The U.S. arm of the exploration company GECO-PRAKLA, part of the Schlumberger Oilfield Services Group, uses a Fujitsu VPX220 single-processor vector machine in its Houston-based computing center for seismic analysis. Interestingly, in spite of press announcements to the contrary [15], we believe that this machine has been purchased in part because it has available an IBM-like operating system which PRAKLA requires to maintain compatibility with other branches of the company worldwide. Supercomputers manufactured by Japanese companies are the only ones with this characteristic. Two other Japanese supercomputers in the U.S. used predominantly for petroleum-related science are the Fujitsu VPX system at TimeSlice Technology, Inc. (see Service Sector below) and the NEC SX-3/22 at the Houston Advanced Research Consortium (HARC) [16]. HARC is a non-profit joint venture among several organizations, mostly universities in Texas and Louisiana.

The financial sector is one of the most interesting and fastest growing users of high-performance computing. In the past, business computing and scientific computing have been seen as two completely separate computer markets with non-overlapping requirements and goals. Although true coalescence of the two domains is still a long way off, today we are beginning to see recognition that massively-parallel processor (MPP) supercomputers designed originally for scientific computing may be an acceptable

alternative to commercial mainframe data center computers [17]. Part of this "recognition" comes from MPP supercomputer vendors realizing that their very survival may depend on their ability to sell machines in the commercial sector, given (a) the larger market for "business" computers, and (b) the decreasing market for large machines in the scientific market.

One estimate puts the potential for supercomputing in the financial services industry at about five times that of the automotive industry [18]. Analysts have also noted that beyond the number of potential customers, the problems requiring computational assistance in the financial industry will have a much larger overall impact on the economy than the design simulations performed in manufacturing industries; and they have the potential for much larger payoffs relative to the price of the supercomputer [19]. In the use of a high-performance computing machine to get results more quickly than the competition, the goals of financial services are the same as most other supercomputer users. Here the result might be a forecast of adjustable mortgage rates rather than, say, a car fender.

There are two basic kinds applications within the financial sector. The first involves attempting to apply sophisticated mathematical models to predict the behavior of securities and futures investments [18, 19]. Japanese firms were early subscribers to this idea (see below) but recently American institutions such as Merrill Lynch, New York's Citibank, and the Federal Home Loan Mortgage Corp. (Freddie Mac) purchased small Cray systems for financial modeling. Several large banks are using scientific workstations for similar tasks, and Prudential Securities' Financial Strategies Group is using a massively-parallel system (described separately below). However, not all of this kind of modeling needs to be done using traditional "number-crunching" techniques on supercomputers. Instead, new, more esoteric solutions, such as neural computing are being investigated, and have met with some success [20].

The second application for high-performance computers in the financial sector is one that has arisen very recently, and involves such areas as fraud detection and "database mining" (using databases of customer information to discern trends that will increase profitability) [21]. Strictly speaking, these do not fit within the traditional definition of supercomputing given in Chapter 1, because they do not require highly-accurate numerical results, nor do they necessarily involve large amounts of computing for each datum, an important characteristic of conventional scientific computing. However, enormous quantities of data are involved, and this is why massive parallelism is seen as playing an important role.

Several U.S. companies with interests in the financial services sector have been quick to see the advantage of parallel machines in their work. Prudential Securities' Financial Strategies Group was the first, using an Intel system starting in 1989 as a high-performance addition to DEC VAX equipment. Now American Express and Dow Jones News Retrieval Service also use parallel systems, although more for data management than for modeling. The recent demise of the MPP vendor Kendall Square Research (KSR) may have a negative impact on the development of financial and/or large-scale data-management applications on massively parallel systems. KSR had formed alliances with several companies, including Electronic Data Systems Inc. (EDS), AMR Corp. (the parent company of American Airlines) and others, to

enhance data processing capabilities using KSR's MPP computer [17]. The status of these alliances is not known at this time.

A case study of MPP usage at Prudential Bache has been published [22]. The key to Prudential's early success in parallel computing was that their Intel parallel system was incorporated into an existing mainframe network, thereby providing improved performance with no change in the user interface. It is interesting that even though Prudential's Intel parallel system was certainly low-performance by today's standards, Prudential was still able to accomplish its goals using the system. A financial model that had previously run in ten minutes could be run in seconds, allowing a broker to run it while on the telephone with a customer. Recently Prudential upgraded their MPP to Intel's latest system, called the "Intel Paragon." Early in its history, the Paragon was believed to have suffered from poor hardware and software reliability [23], and since system stability is one of the foremost requirements for financial computing, it will be interesting to see if Prudential continues to be satisfied with the Intel MPP.

We wonder why more financial firms do not use supercomputers. The reason may be that much of the difficulty in financial forecasting lies in the development of the computer model (i.e. deciding what factors are to be included and how to actually do the prediction), rather than a lack of computer "horsepower." Perhaps, though, there is a useful analogy with weather prediction, which, although also subject to numerous inherent uncertainties in model development and parameter selection, has been a staunch supporter of supercomputing all along. The difference, of course, is that weather modeling is generally supported by government (e.g. NCAR, NOAA, Navy POPS, etc.), whereas financial computing is largely in the private sector.

Finally, we note that financial services computing is one of the key areas in which collaborations between U.S. companies, federal agencies, and DOE National Laboratories is taking place. At Los Alamos, topics involving neural networks, fraud detection, and database manipulation are being jointly investigated with entities such as Citibank and the U.S. Department of the Treasury, even though neither of these customers currently uses supercomputers in their work.

Supercomputers play a critical role in automobile industry R&D. Recently supercomputers have been seen as a primary means of improving product quality to challenge Japanese competition [24]. Within automobile companies, supercomputers are used to guide, to focus, and to reduce laboratory experimentation, thereby reducing design time. One automobile engineer has noted the importance of the supercomputer in allowing previously-untenable three-dimensional model simulations to be done [25].

Among automobile manufacturers in the U.S. Ford is by far the strongest in terms of supercomputing power, having two of Cray Research's most recent vector machines and one smaller vector system as well. In fact, Ford has the second most powerful computing system installed at any non-vendor commercial entity in the world (after Exxon). Chrysler and General Motors also each have one Cray Research system. Although it is not included in our database, Chrysler is the first of the American

automobile manufacturers to use parallel processing; it will install a Cray Research T3D system sometime soon. Ford also plans to use parallel processing, and will install a system that consists of both Convex and Hewlett-Packard components. More discussion of supercomputing usage within the automobile sector will be presented in the next chapter.

Of the chemical/pharmaceutical companies using supercomputers, one, DuPont, is an integrated chemical maker and the other five are pharmaceutical companies. The usage of supercomputing at DuPont is well-documented [26-29], and consists of (a) process optimization for chemical plants and (b) computational chemistry to support research, in roughly equal amounts. DuPont is unique among all commercial high-performance computing users in that it has estimated rather precisely the actual dollar value returned by its supercomputing efforts. Their estimate is that increases in product yields and decreases in manufacturing costs obtained through simulation now save the company about \$250 million per year.

DuPont's computational needs would optimally cover a wide range of computing platforms, of which large, expensive, vector supercomputers are only one. For example, a significant portion of computational chemistry simulations do not vectorize and thus may not be optimally suited for vector machines. However, DuPont stresses the need to do some of the process simulations in near real time, so that optimizations can be made at the plant as new situations (such as purity level of an ingredient) arise; for this kind of work, the rapid response time of a vector machine is critical. Additionally, DuPont would like to move into the massively-parallel arena, but feels that the software technology for such is not currently mature enough. The software used for process simulations on its Cray Research supercomputer was purchased from third-party vendors [27], and is not yet available for machines such as the T3D.

The five U.S. pharmaceutical companies in our database are among the largest in the United States. Merck, Bristol-Myers-Squibb, Eli Lilly, Marion Merrill Dow, and Pfizer all use supercomputers for quantum chemistry and molecular modeling of potential pharmaceutical agents [30]. Process modeling is also important at pharmaceutical companies. Eli Lilly uses an older Cray Research CRAY-2 system and will may upgrade to a new machine, quite possibly massively parallel, during 1994.

Several other pharmaceutical companies worldwide that are using true supercomputers in their R&D include Schering and Bayer in Germany [31], Taisho in Japan (which has a small desktop system and is therefore not included in our database), and two other companies in Japan that have old, smaller Fujitsu or NEC supercomputers and are also not included in our list. Although Cray Research and other external observers believe that usage within the pharmaceutical industry will increase, there are signs that the opposite may be true, at least in parts of Europe. For example, we learned recently that none of the major pharmaceutical manufacturers in Switzerland (Ciba-Geigy, Roche, and Sandoz) have much interest in supercomputing [32]. Sandoz is content to use Digital Equipment Company VAX mainframe systems and the other two use scientific workstations. Sandoz at one time used an Alliant system [33], an early parallel processing computer made by the now-defunct company Alliant Computer Systems.

The service sector includes companies that sell supercomputer cycles and/or offer value-added services such as applications, research, or consulting. It is becoming more difficult to identify precisely which companies belong in this sector because recently, companies with primary business interests in other areas have begun selling cycles or offering services on machines that they originally purchased for their own computing needs. This is happening because of two primary reasons. As an example of the first, consider that Cray Research recently purchased MSCI, the company that runs supercomputers for the University of Minnesota. We suppose that it makes sense for a supercomputer vendor like CRI to enter the services business. Many computer vendors are finding it difficult to survive through hardware sales alone, and supercomputers vendors are particularly vulnerable. Second, cutbacks in various business sectors have resulted in a surplus of supercomputing capacity. Consider that Lockheed Information Technology Corporation, which is the supercomputer data services arm of the Lockheed group of companies, recently began offering user services to outside customers in order to reduce supercomputing costs to Lockheed and to make use of unused cycles on their two Cray Y-MP supercomputers. However, surplus supercomputing capacity has also caused two early computer service firms, Boeing Computer Services and Grumman Computer Services to stop selling time on their machines.

Thus, although many companies with supercomputers could be classified as being in the service sector, we include just two in our database, TimeSlice Technology, Inc. of Houston, Texas and Booz, Allen, & Hamilton, located in McLean, Virginia. TimeSlice leases a Fujitsu VPX260 for its value-added supercomputer out-sourcing business, in which applications, research, and consulting services are offered to companies in a variety of business sectors [34]. The VPX260, one of three Fujitsu systems installed in the U.S., is only a single-processor vector supercomputer but it is the most powerful single-processor system installed in the United States. Booz, Allen & Hamilton is one of the nation's largest consulting companies. They use a small version of the Connection Machine CM-5 to run neural net programs related to a variety of both military and commercial applications, few of which are traditional numerical scientific simulations. Apparently, Booz, Allen & Hamilton management purchased the CM-5 anticipating future use for it, after determining that massively-parallel computing was an important advanced technology in which the company should be involved.

2.4.3 Supercomputing Usage by Sector in Japan

Figure 23 shows the distribution of supercomputer installations in Japan from the 580-supercomputer database and Figure 24 shows the distribution of computing power according to LINPACK. Figures 25 and 26 show the sector distribution of vector supercomputer installations and vector supercomputing power in Japan, and in Figures 27 and 28 the corresponding sector distribution for parallel supercomputers is shown.

The sector distribution of supercomputer installations in Japan is radically different from that of the United States. (Compare Figures 27 and 28 with Figures 17 and 18 on page 27.) And because of the NWT machine and six Fujitsu VPP500s installed in Japan, the difference between the sector distribution

according to number of installations and according to power is considerably more distorted than in the United States. In particular, there is an enormously unequal distribution of parallel supercomputing power in Japan, with the two government categories, non-defense and aerospace, accounting for fully one-half of all parallel installations and 80% of parallel computing power.

According to the author, the United States is a country of immigrants. The author states that the United States is a country of immigrants, and that the United States is a country of immigrants. The author states that the United States is a country of immigrants, and that the United States is a country of immigrants.

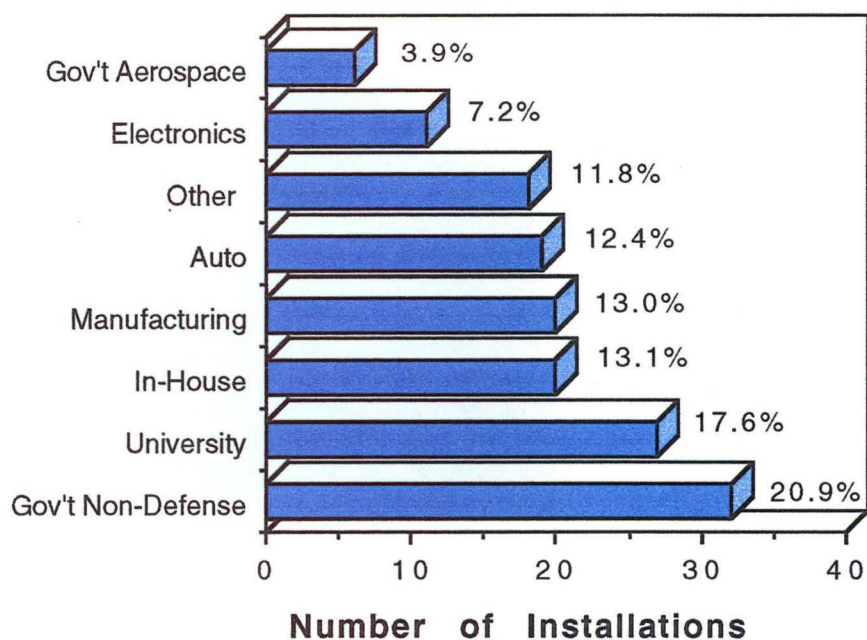


Figure 23. Sector Distribution of Supercomputer Installations in Japan (155 systems).

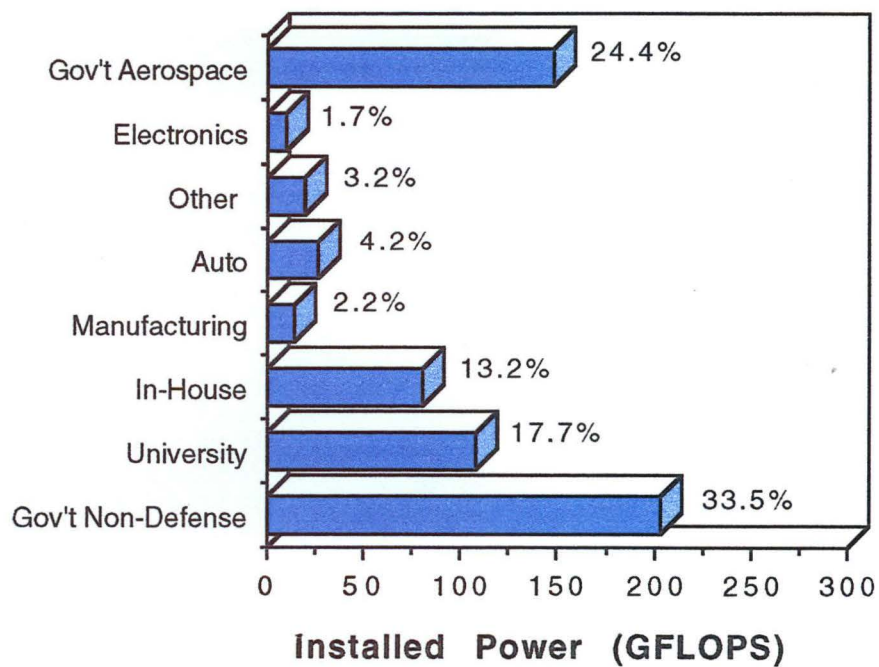


Figure 24. Sector Distribution of Supercomputing Power in Japan.

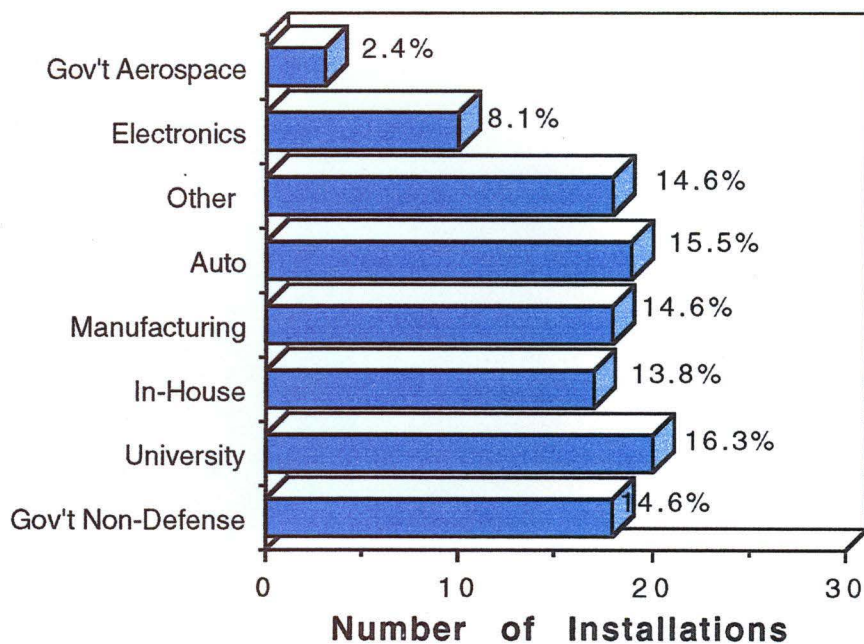


Figure 25. Sector Distribution of Vector Supercomputer Installations in Japan (124 systems).

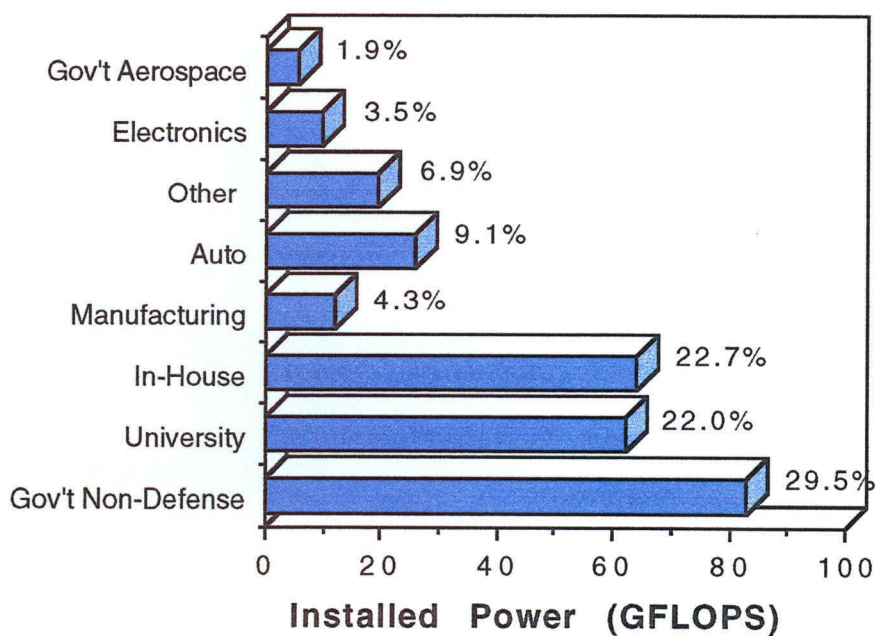


Figure 26. Sector Distribution of Vector Supercomputer Power in Japan.

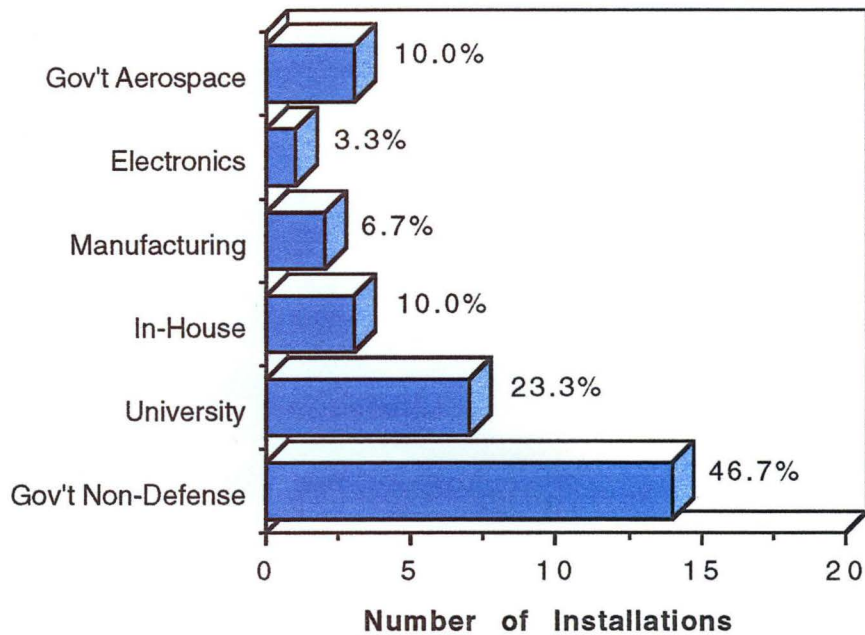


Figure 27. Sector Distribution of Parallel Supercomputer Installations in Japan (31 systems).

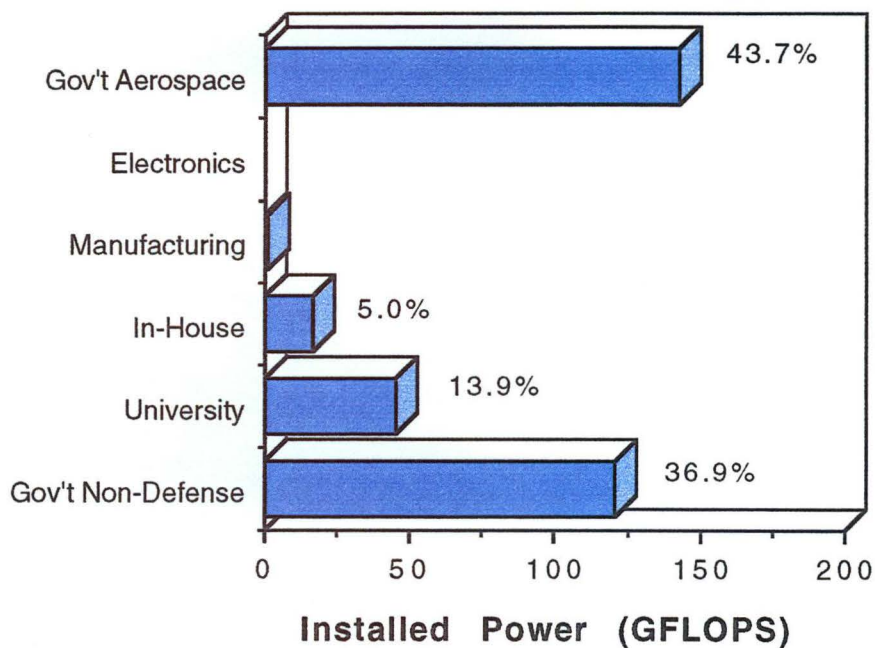


Figure 28. Sector Distribution of Parallel Supercomputer Power in Japan.

2.4.3.1 Japanese Government Usage

In Japan, the largest sector in terms of both number of installations and power is the government non-defense sector. It accounts for about 21% of all installations and about one-third of installed power in Japan.

Non-defense government research labs in Japan have almost three times as much supercomputing power as non-defense government labs in the United States. It is true that large parallel systems such as the Fujitsu VPP500 and Cray Research T3D can, perhaps, artificially inflate this ratio, and there are four of these machines in the Japanese government laboratories. However, the Japanese government labs also maintain a four-fold *vector* supercomputing power advantage over the American government non-defense laboratories.

The areas of research in the Japanese government non-defense category encompass a wide range of interests, as shown in Table 2.3. The main difference between these and the non-defense U.S. government laboratories is that (1) there are more of them in Japan, and (2) there are no multi-disciplinary laboratories analogous to Argonne National Laboratory or Oak Ridge National Laboratory in Japan; all government labs have a more restricted scientific charter. Many of the Japanese laboratories have no direct U.S. Government counterpart but some of the work done in them is similar to that done at U.S. companies or universities.

The Japanese government non-defense laboratories have taken the lead in Japan in using parallel processing. One important example is the Advanced Telecommunications Research Institute (ATR), which is experimenting with several kinds of parallel machines, all from American vendors. In the past, Japanese research in high-performance computing has characteristically placed more emphasis on non-numerical computation than much of the research in the United States. Although parallel processing research in Japan is still at a relatively early stage, some of it is already oriented towards non-numerical computing. For example, much of the research at ATR consists of studies in cognitive research, speech processing, and computer vision [35]. Also, the Real World Computing Program will use a CM-5, an Intel Paragon, and a machine it will develop on its own, for development of information processing technologies that are intended to mimic the way humans process data.

However, more traditional scientific simulations are also under way on parallel machines installed at places like the Angstrom Technology Partnership (ATP) and at RIKEN, the Institute for Physical and Chemical Research. Our visit to ATP, described in the next chapter, suggests that research at Japanese labs like ATP will concentrate much more on pure physical science rather than on the variety of issues associated with parallel processing, such as languages, architecture, and new algorithm development. We believe this is another important difference between the highly-focused Japanese labs and the more inter-disciplinary environments at U.S. National Laboratories, which are much more involved in the

Table 2.3: Government Non-Defense Supercomputing Sites in Japan

Site	Vendor	Model	Type	LINPACK Rating (MFLOPS)
Advanced Telecommunications Research Institute	Intel	iPSC/860/6	P	0.4
	KSR	KSR1-96	P	2.5
	TMC	CM-2	P	1.3
	TMC	CM-5	P	3.8
Agency of Industrial Science and Technology	TMC	CM-5	P	3.8
AIST Research Information Processing System	Cray	C90/16	V	13.7
Angstrom Tech. Partnership	Fujitsu	VPP500/30	P	33.0
	TMC	CM5E/128	P	12.8
Communications Res. Lab.	Fujitsu	VPP500/10	P	11.5
Disaster Prevention Res. Inst.	Cray	Y-MP2	V	0.6
Inst. For Computational Fluid Dynamics	TMC	CM-2	P	1.3
Inst. for Physical and Chemical Res. (RIKEN)	Fujitsu	VPP500/28	P	30.8
Institute for Fluid Science	Cray	Y-MP	V	2.1
Institute for Molecular Science	Hitachi	S820	V	1.8
	NEC	SX-3/34r	V	17.4
	Fujitsu	VP2600	V	4.0
Japan Atomic Energy Research Institute	Fujitsu	VP2600	V	4.0
	NEC	SX-3/41R	V	5.8
	MasPar	2216/16k	P	1.6
Nat'l Cancer Res. Center. Inst.	MasPar	2216/16k	P	1.6
	NEC	SX-3/14	V	5.0
Nat'l Inst. for Environmental Science	NEC	SX-3/24R	V	11.6
Nat'l Inst. for Fusion Science	NEC	SX-3/24R	V	11.6
Nat'l Lab. for High-Energy Physics	Hitachi	S820	V	1.8
Nat'l Research Inst. for Metals	Hitachi	S3800/3	V	5.3
Nat'l Space Observatory	Fujitsu	VP200	V	0.4
National Astronomical Lab.	Fujitsu	VP200	V	0.4
National Fusion Research Lab.	Fujitsu	VP200	V	0.4
Power Reactor and Nuclear Fuel Development Corp.	Cray	T3D/128	P	10.7
	Fujitsu	VP2600	V	4.0
	Fujitsu	VP2600	V	4.0
Real World Computing Program	Intel	Paragon	P	2.0
	TMC	CM-5	P	3.8

P = Parallel

V = Vector

latter areas. It should be pointed out that this is not necessarily an unqualified advantage for Japan. The entire field of parallel processing is so new that few universally-recognized standards exist. By researching the nascent topics in this area U.S. researchers are able to help establish or influence standards, thus ensuring that their respective requirements are met.

Another interesting aspect of the Japanese Government laboratories is that even though some of them are very well endowed with supercomputer hardware, there are often few users for the systems, at least

many fewer than would be demanding time at comparable U.S. facilities. Three examples of this include our own experience at ATP, where we saw the queue for the Fujitsu VPP500 with only a handful of jobs awaiting execution, the new 24-GFLOP Hitachi S3800 at the Institute for Metals Research, which had only "four or five" users in its first two months [36], and a statement made by a U.S. NASA researcher who noted that because there are fewer active CFD researchers in Japan and all seem to have access to many supercomputers, more supercomputing time is available per researcher than in the United States [37].

Finally, we note that one of the more important users listed in Table 3 is the Power Reactor and Nuclear Fuel Development Corporation, known in Japan as PNC. This is a quasi-government organization that is developing Japan's *Monju* prototype fast breeder reactor. PNC is well-endowed with supercomputer power, with two high-end Fujitsu vector machines and a new CRAY T3D massively-parallel system; compare, for example, PNC's 18-GFLOPS total power with the various organizations in Table 2.1. Presumably PNC's supercomputers are used for elementary particle and radiation transport simulations, which can be used to predict the reactor flux and help design shielding for the reactor facility, and for simulation of coolant-loss accidents. Interestingly, one code in use at PNC for radiation transport may very well be the MCNP (Monte-Carlo Neutron Photon) code developed originally at Los Alamos National Laboratory. (We know that MCNP is used at other institutions in Japan such as Japan Atomic Energy Research Institute (JAERI)). Recently, we have also learned that PNC is attempting to carry out computer simulations of the extraction process by which radioactive waste products are separated from the bulk waste material [38].

There are only two supercomputers in Japan considered to be in the defense sector (a Hitachi S820, an old machine with little power by today's standards, and a small CRAY Y-MP system). We know little about these two sites.

2.4.3.1.1 Government Aerospace Research in Japan

The increasing emphasis Japan is placing on its small but rapidly growing aerospace industry is significant, although perhaps not extensively known in the West [39]. The computational resources associated with Japan's national aerospace laboratories (there are two: the National Aerospace Laboratory and the Institute for Space and Astronautical Science, both in Tokyo) underscore the extent to which computer simulation is expected to play a role in the development of this industry. Because of the NWT, Japanese Government-funded aerospace laboratories have about twice the total computing power of NASA. That is, 17 NASA machines have one-half the power of six Japanese systems. The approach of the Japanese Government is to emphasize one or two very powerful space research computational facilities, rather than a variety of smaller, less capable establishments, as NASA has done. More importantly, by funding the design and the construction of the NWT, they are also emphasizing a new, state-of-the-art, very expensive, and to a certain extent, heretofore unproven technology. Clearly, such a strategy is designed for the purpose of generating breakthroughs rather than incremental improvements

in aerospace knowledge; perhaps the NWT should be viewed as the computational equivalent of the U.S. Apollo program.

2.4.3.2 University Supercomputing in Japan

Japan has only about one-third as many supercomputers in its universities as the U.S. does. The combined statistics further suggest that Japanese universities have about 37% of the computing power of U.S. universities but the situation is actually worse than those numbers suggest. This is because there is one Fujitsu VPP500 system installed at a Japanese university (University of Tsukuba) and it alone accounts for 31% of all Japanese university computing power. Much of the remaining university computing power, which is marginal compared with that of the U.S., is also concentrated at a few larger schools, such as Tohoku and Tokyo Universities.

Interest in parallel processing in Japanese universities has picked up slightly in the last year or so. There are now three small Thinking Machines Corp. massively-parallel CM-5s at Japanese universities, including one of the latest models, called the CM-5E. However, these are all small configurations of the machines, and many Japanese researchers still prefer vector supercomputing for their research. One researcher at a Japanese Government research lab recently referred to the 64-processor Thinking Machines CM-5 he had run on at Hokuriku Graduate University (run by Japan's Agency for Industry and Science and Technology, AIST) as a "toy," and therefore he decided to purchase a 3-processor Hitachi vector machine for his university research institute [36].

Because they have only small configurations of parallel machines, there is still more vector processing power at Japanese universities than there is parallel processing power, a situation which is quite the opposite of the U.S. universities. In fact, Japanese universities have more vector processing power than do U.S. universities. This is in spite of the fact that there are more vector machines installed in U.S. universities than there are at Japanese universities. Many of the vector machines installed at Japanese universities are the most powerful configurations in their respective vendor families. For example, Tohoku University, a very well-endowed university in the northern part of Japan, has a four-processor NEC SX-3/44R, currently the flagship of the NEC vector line. Similarly, Tokyo University has a four-processor Hitachi S3800, which is the most powerful machine Hitachi makes; in fact, it is the most powerful vector supercomputer made by any vendor. Another interesting fact about these two machines is that they are currently the only fully-configured versions of these machines in Japan installed outside the companies that make them. In order to find the most powerful Japanese-made vector processors in Japan, one must go to a university, and not to a government research laboratory. (Recently, a four-processor NEC SX-3R supercomputer was installed at the Canadian climate modeling facility in Dorval, Ontario, and this is the only fully-configured, multiprocessor Japanese supercomputer installed outside of Japan.)

In summary, the contrast between Japanese and American universities seems to be as follows. There are more American universities with supercomputer installations, making both vector and parallel

processing equipment available to a larger number of students and researchers than are Japanese universities. This is especially true for parallel systems, and it has the effect of enabling a much larger class of users to gain supercomputing experience. In both the U.S. and Japan, a few selected universities have accumulated large concentrations of supercomputing power, thereby emphasizing the research capabilities of a small class of investigators. In the U.S. this concentration is largely of parallel machines and in Japan it is largely of vector machines. Because parallel machines are able to accommodate scientific problems of a larger scale, U.S. university researchers may have an advantage over most of their Japanese counterparts.

2.4.3.3 Commercial Supercomputing Usage in Japan

2.4.3.3.1 In-House Supercomputing in Japan

The publicly-available data suggest that there is roughly an equivalent amount of supercomputing power used in-house in Japan and the United States, except that in Japan it is 80% vector / 20% parallel, while in the U.S. it is 40% vector / 60% parallel. Recall that in the U.S., of all the commercial supercomputer users, the supercomputer vendors themselves owned the largest number of machines. This is also true in Japan, and in Japan the manufacturers account for an even larger share of the installed commercial supercomputing power.

We believe that our estimate of the U.S. in-house resources is considerably more accurate than is our estimate of the three Japanese manufacturers. We know that there are probably three NEC vector supercomputers used within various branches of the company, but we do not know how many in-house installations of the NEC parallel product, called "CENJU" there are. NEC has a computational science research group that studies physical and chemical properties of semiconductors, new materials, and biomolecules [40]. This is exploratory work, believed to be outside the scope of current NEC business but possibly relevant to future products. A small benefit to NEC's current product line was realized because quantum chemistry software developed by the research is now sold as a product for outside users.

Regarding Fujitsu, we believe that there are a few in-house installations of its VPP500 system. All of these have relatively small configurations, the largest with about 11 processors. One important aspect of in-house supercomputing capabilities in Japan is Fujitsu's establishment of an in-house parallel computing research facility at which several versions of its AP-1000 parallel computer are located and made available to users worldwide via the Internet. Fujitsu has reported that nearly 800 non-Fujitsu users are validated on these systems [41]. Interestingly, Fujitsu's data show that almost all of the Japanese users are from universities and National Laboratories; i.e., there are no commercial users. Because the AP-1000 is not regarded as being a state-of-the-art system in terms of its performance [8], it is doubtful that "break-through" kinds of research are being carried out on it. However, this kind of facility is rare in Japan and the experience in parallel processing garnered by its users is no doubt highly valuable.

Finally, Hitachi has a large number of its supercomputers, mostly older models, installed in-house. Hitachi's business interests, however, are so much larger and more diverse than that of all the other supercomputer manufacturers, that a direct comparison of all its in-house resources is not appropriate. Therefore, for Hitachi's in-house computing, we only counted those machines (eight in all) installed in Hitachi's General Purpose Computer Division. One of these machines may have recently been replaced by the newest Hitachi supercomputer, the S3800, but we have been unable to confirm this.

2.4.3.3.2 Supercomputing in Other Japanese Industries

Eliminating those machines used in-house and those used by government and universities leaves 67 machines installed within non-vendor industries, which is 44% of all Japanese supercomputers. This is the same number of installations as within non-vendor industries in the United States, but in the U.S. the percentage is about 25% of all machines. In both countries machines in the commercial sectors account for a disproportionately small share of total supercomputing power, about 11% in Japan and about 14% in the United States. However, in absolute terms, the U.S. has more than twice as much supercomputing power installed in its non-vendor industries as Japan does.

One-half of all Japanese vector supercomputers are located in Japanese companies. However, Figures 25 and 26 show that the one-half of Japanese vector supercomputers installed in companies account for only about 23% of total vector supercomputing power in Japan, suggesting that many of these machines are lower-end models with relatively low performance.

Hardly any parallel supercomputing power exists within Japanese industries, whereas there is an abundance in U.S. industries. Only 4 of the 31 parallel machines in Japan are installed in companies, compared with 32 parallel systems used by non-vendor companies in the United States. Furthermore, the parallel machines used in Japanese industries are all low-end machines, so that the combined parallel computing power associated with Japanese companies is less than 10% of what it is in the United States. The machine that so dominates the Japanese parallel processing power estimates, the Fujitsu VPP500, has not yet made a sale to any Japanese company; it has only been installed at universities and government labs.

Table 2.4 lists industry sectors in Japan that use supercomputers. The overall comparison between U.S. and Japanese commercial supercomputing may be divided into three groups. The first group consists of commercial sectors in which both U.S. and Japanese companies have supercomputers; these are the automobile, electronics, financial, chemical, manufacturing, and service sectors. For some of these, comparison between the U.S. and Japan is still difficult because in one country or the other, only a small

Table 2.4: Commercial Supercomputing in Japan.

Sector	Number of Machines		Combined LINPACK Rating (GFLOPS)
	Vector	Parallel	
Automobile	18	0	25.3
Chemical / Pharmaceutical	3	0	1.1
Construction	6	0	9.3
Electronics / Telecommunications	10	1	10.2
Financial	3	0	2.2
In-House	17	3	80.1
Manufacturing	17	2	12.9
Service	1	0	1.7
Steel	3	0	2.6
Transportation	1	0	1.2
Utility	1	1	2.2
Total	80	7	148.7

number of machines are involved. For example, in the service sector, we have listed only two machines in the U.S. and one in Japan.

The second group is made up of four sectors in which Japanese companies use supercomputers but in which no supercomputers are used in analogous U.S. companies (construction, steel, transportation, and utility). The third group consists of a single sector in which there is supercomputing usage in the U.S. but not in Japan, namely the petroleum industry. Note, also, that in the U.S., we have the commercial aerospace sector, while in Japan there are no companies using supercomputers whose business is entirely restricted to aerospace. Japanese companies with interests in aerospace, such as Mitsubishi Heavy Industries, are listed in the manufacturing sector.

We begin by considering sectors in which both the U.S. and Japan use supercomputers. The largest non-vendor industrial user of supercomputers in Japan is the automobile industry. In the U.S., automobile companies are the second largest sector, after petroleum. Table 2.5 compares supercomputer resources within the automobile industry in both the U.S. and Japan. (Fuji Heavy Industries, Inc. (FHI), is a manufacturer of minivans and four-wheel-drive vehicles.) The 12 automobile companies in Japan have a total of 18 supercomputers and the 3 companies in the United States have a total of 8. However, supercomputing power in the U.S. auto companies is actually greater than it is in Japan, largely because of the two big Cray Research C90 systems at Ford.

Table 2.5: Supercomputing in the Automobile Industry in the United States and Japan.

Company	Machine	LINPACK GFLOPS
Ford	CRAY C90/16; CRAY C90/8; CRAY M98/6; Convex S3880	13.7; 6.8; 1.4; 7.9
General Motors	CRAY M98/4; CRAY Y-MP/3	1.1; 0.8
Chrysler	CRAY C90/4; CRAY Y-MP/6	3.2; 1.6
Total for U.S.:	8 systems	36.5
Daihatsu	CRAY Y-MP/2	0.6
Fuji Heavy Industries	Fujitsu VP2600	4.0
Hino Motor	NEC SX-31/L; CRAY Y-MP/2	0.6; 0.6
Honda	CRAY Y-MP/3	0.8
Isuzu	CRAY Y-MP/2	0.6
Mazda	CRAY Y-MP/2; CRAY Y-MP/2; NEC SX-2	0.6; 0.6; 0.8
Mitsubishi Motors	CRAY C90/2; CRAY Y-MP/4	1.7; 1.1
Nissan	CRAY Y-MP/6	1.6
Suzuki	Hitachi S3800	1.7
Toyota	CRAY Y-MP/6; Fujitsu VP2200	1.6; 0.8
Toyota Auto Body	Fujitsu VP2200; Fujitsu VP2100	0.8; 0.6
Toyota Central Research & Development Lab	NEC SX-3/14	5.8
Total for Japan:	18 systems	24.9

Three automobile companies in Japan use supercomputers from more than one vendor in their R&D. In each case, they have a Cray Research system and a Japanese-manufactured system. One of these companies is profiled in some detail in the next chapter, but the reason for having two kinds of machines is the same at all three companies. The Cray systems are used for running crash simulations and performing structural analyses, because the third-party software for doing these is superior to that which is available on Japanese-manufactured supercomputers [42, 43]. The Japanese machines, a Fujitsu VP2200 at Toyota, an NEC SX-2 at Mazda, and an NEC SX-3/1L at Hino, are used for computational fluid dynamics (CFD) simulations using in-house developed software. CFD simulations help define body shape, reduce drag, measure heat loss from the engine, and, in a more specialized form, simulate combustion in the engine.

In contrast with structural and crash studies, there is still much debate about the usefulness of fluid simulations in automobile design [44, 45]. The debate arises from potential inaccuracies in the mathematical models used and from the fact that simulations still require too much computational time to be useful on a day-to-day basis. This is an area in which advanced architecture computers will make a significant contribution (if companies are willing to pay high prices for them).

In contrast, one of the most interesting aspects of supercomputing within the Japanese automobile industry is that there are two successful companies that do not use Cray supercomputers at all. Cray machines, by virtue of their superior crash simulation software, are used by virtually all other automobile companies in Japan and the U.S., as well as the overwhelming majority of the important automobile companies in Europe and South Korea, including Daimler-Benz, Fiat, Renault, Citroen, Opel, Volkswagen, Audi, BMW, Kia, and Hyundai. Many of these companies have more than one Cray Research system; and often, one of the two is dedicated entirely to crash analysis. Yet, Japan's Suzuki and Fuji Heavy Industries use only Japanese-manufactured supercomputers, probably with in-house crash simulation software. (These two companies would have made excellent candidates for case studies but we were unable to arrange a visit.)

The manufacturing sector is the second largest commercial, non-vendor user of supercomputers in Japan. In contrast, in the U.S. it is one of the smallest, with only a single company, Westinghouse, represented [46]. The companies in the Japanese manufacturing sector that use supercomputers are listed in Table 2.6. Although there are a variety of machines used by these companies, the single most common characteristic is relatively low performance. There are two parallel machines in this list, (both installed at Mitsubishi-group companies) and even these are low-performance machines (in fact, almost ridiculously small). Such machines may be used for experimentation with parallel processing technologies with the view that, if successful, larger and more powerful machines would be purchased in the future.

The companies in Table 2.6 also represent a wide range of product types, both between companies and sometimes even within a single company. In fact, categorization of Japanese companies according to our scheme is difficult because of the wide range of product types many span. In particular, the distinction between the "manufacturing" and "electronics" categories is somewhat arbitrary. For example, Mitsubishi is involved with both nuclear power as well as production of DRAM memory. And, as noted above, several of these companies, notably Hitachi, Mitsubishi Heavy Industries, and Kawasaki Heavy Industries, are involved in production of defense-related products for Japan, such as aircraft and space systems.

Table 2.6: Supercomputers in Use in the Japanese Manufacturing Sector.

Site	Business	Machine	LINPACK GFLOPS
Mitsubishi Electric	Heavy Machinery / Electronics	NCube nCube2	0.9
Hitachi Adv. Res. Lab	Heavy Machinery / Electronics	Hitachi S820	0.9
Hitachi Central Res. Lab	Heavy Machinery / Electronics	Hitachi S820	0.9
Hitachi	Heavy Machinery / Electronics	Hitachi S820	0.9
Hitachi	Heavy Machinery / Electronics	Hitachi S820	0.9
Mitsubishi Heavy Industries	Heavy Machinery / Electronics	CRAY Y-MP/2	0.6
Hitachi Zosen Corp.	Shipbuilding	Hitachi S3600	0.8
Toyo Tire	Auto Tires	Hitachi S3600	0.8
Kawasaki Heavy Indust.	Heavy Machinery / Aerospace	Fujitsu VP2200	0.8
Sumitomo Rubber	Auto Tires	NEC SX-3/1LR	0.8
Komatsu Mfg.	Heavy Machinery	NEC SX-3/1L	0.7
Nippondenso	Automobile Parts	CRAY Y-MP2	0.6
Mitsubishi Electric	Heavy Machinery / Electronics	Intel iPSC/860-16	0.4
Yamaha Motor	Motorcycles	CRAY Y-MP/1	0.3

Because Japanese manufacturing companies have such a wide range of products, their general-purpose vector supercomputers support a wide range of simulation types. As an example, consider the use of the Cray Research system at Mitsubishi Heavy Industries (MHI). This company uses a 2-processor Cray system with very large memory at its Kobe Shipyard and Machinery Works for structural analysis, fluid dynamics, electromagnetics, molecular dynamics, and graphics visualization [47]. Examples of simulations MHI has carried out related to its products are safety analysis of spent nuclear fuel shipping casks, metal forming analysis of steam generator channel heads, grounding analysis of a ship bottom, flow analysis within nuclear reactor vessels, boiler furnaces and diesel engines, and structural analysis of tunnel excavator heads.

MHI is an interesting example of a Japanese company that began its supercomputing efforts as a result of a recent (1989) management policy that sought to significantly modernize its design and development activities. (It is not the only Japanese company that has reported bringing in a supercomputer as part of a

modernization effort; see the sections below on the construction, electronics, and steel industries.) MHI initially attempted what it describes as its "revolutionary engineering work" using a general-purpose IBM mainframe system. Users of this system were charged a fee based on their usage, but because the machine was so slow, MHI decided to introduce a small CRAY X-MP/EA system in order to reduce usage fees. MHI estimated that the low-end Cray system performed structural computations at 10% of the cost of the IBM machine, and later upgraded to a more powerful Cray Research system when the smaller, older machine became saturated with user jobs. MHI's experience would seem to be one that argues clearly for the economic efficiency afforded by supercomputing (relative to mainframes) in spite of the initial cost of the machine.

Two automobile tire manufacturers in Japan (Sumitomo Rubber and Toyo Tire) use supercomputers [48]. We are aware of two other tire manufacturers (Continental A.G. and Michelin [49], both in Europe) where supercomputers are also used. The European companies use small Cray systems (they are below the performance threshold for our database) but the Japanese manufacturers use vector machines from NEC and Hitachi. In tire design, numerical modeling of tire behavior under varying loads and obstacle environments can be performed as a function of changes in tire geometry, tire and belt material, and tensions in the reinforcing cords. Previously, such data were obtainable only from physical laboratory or test track experiments, and so the use of the supercomputer clearly allows shorter development cycles and reduced prototype costs. This is true even though tire simulation presents numerous difficulties not encountered in other types of structural modeling such as cars, because the composite materials used in tires deform much more readily. Thus, not only are more complex mathematical models required, but so are non-uniform computational grids which tend to lengthen the computation time.

The combined electronics/telecommunications sector is the third largest commercial non-vendor user of supercomputers in Japan, with eight vector machines and one small parallel system used in six different companies. In the U.S., there are three vector and four parallel machines in six companies. Combined supercomputing power in Japan's electronics/telecommunications sector is nearly twice that of the United States' electronics/telecommunications sector.

Recall that some of the companies categorized under other sectors also have business interests and perform supercomputer simulations that could be classified as "electronics," such as Hitachi and Mitsubishi Electric. Some of the firms listed in the Electronics group are semiconductor manufacturers and use their supercomputers for integrated circuit (IC) device simulations, either for ICs to be included in their own products (such as Matsushita) or for "chips" to be sold "as is," such as Toshiba. Toshiba, which is the largest DRAM producer in Japan, uses a Cray multiprocessor vector machine for its work and has developed some of the software for it internally [50].

Semiconductor device simulation is an interesting example of an application that cannot be carried out effectively without some sort of high-performance computing engine; however, traditional vector supercomputers have known inadequacies in their deliverable performance on this type of simulation. For this reason, some of the Japanese semiconductor manufacturers have developed their own special-

purpose computing systems dedicated to this task. The machines developed by NEC have been discussed previously [51] and the system developed by Hitachi is covered in the next chapter.

Nippon Telegraph and Telephone (NTT), currently uses its CRAY-2 supercomputer for a variety of tasks in support of research in optics, semiconductor physics, and material science. The CRAY-2 system is expected to be upgraded during early 1995. NTT also has a large distributed computing research group that uses a small Intel iPSC/860 parallel processing system.

Companies in the electronics/telecommunications sector also use supercomputers for simulations related to the non-electronic portions of their consumer products [52]. For example, Sanyo Electric, which has a small NEC SX-3 system, has used the supercomputer to simulate resin flow in the cabinet of an HDTV system [53]. In this case, design efficiency was increased by having the supercomputer optimize molding conditions, producing results that would be difficult to obtain by any other method. Interestingly, Sanyo states that it deliberately introduced the SX-3 computer system in an attempt to promote Computer Aided Engineering (CAE) throughout the company by creating a user environment that is relatively easy for engineers to use. To this end, Sanyo is itself developing much of the application software it needs. An important point is that Sanyo is another case where *management* decided to encourage the use of supercomputing in order to increase efficiency over existing mainframe and workstation capabilities and to encourage greater participation in CAE throughout the company.

Three Japanese chemical companies in our database, Sumitomo Chemical, Asahi Chemical, and Mitsui Toatsu Chemical, use supercomputers. Although all three of these companies are large comprehensive chemical makers like DuPont (Asahi is the largest in Japan), in all respects they are much smaller than DuPont. In our recent trip to Japan we were told that chemical companies in Japan that have supercomputers generally use them only for basic chemical research, rather than for process modeling as DuPont does. As mentioned above, DuPont's ability to simulate its production facilities using Cray systems has yielded enormous financial savings for the company; thus, it seems that Japanese companies are lagging in this important area. Since the process modeling software is available from third-party vendors, it is difficult to understand why Japanese chemical companies do not use it.

We are aware of three companies in the financial sector in Japan that use supercomputers, as opposed to four in the United States. The Japanese companies are Nikko Securities, Daiwa Securities, and Yamaichi Securities and they all use relatively old and lower-end Japanese-manufactured vector machines. There are indications that Cray Research, buoyed by its recent sales to customers such as Merrill Lynch and Federal Home Mortgage Corporation (Freddie Mac), is expecting additional sales in Japan. Whereas several financial companies in the U.S. use large parallel systems (American Express and Prudential Securities), only vector machines are used in Japanese securities firms. Because of this, there is currently much more supercomputing power applied within the financial sector in the U.S. than there is in Japan.

The remaining sector in which both the U.S. and Japan have supercomputing activity is the service sector. In Japan we are aware of two companies whose primary line of business involves supercomputing

services. (One of them is actually not in our database because they no longer use true supercomputing equipment, having substituted workstations for a CRAY X-MP). Both companies, Computer Technology Integrator, Ltd. (CTI) and CRC Research Institute, are profiled in the next chapter. Although both CTI and CRC accept customers from all applications fields, the majority of their work is concentrated in selected areas: Electric power and aerospace for CTI and nuclear power for CRC. (The same is basically true for one of the two service companies in the U.S., (TimeSlice), which concentrates largely on petroleum industry customers.) Another (in)famous but no longer active supercomputer service company in Japan is Recruit, Ltd., which, in its heyday, owned several Cray Research supercomputers and an NEC SX-2, and also maintained its own Institute for Supercomputing Research. Various misfortunes befell the company causing its present owner to opt out of the supercomputer out-sourcing business, although Recruit still maintains some activity as a third-party supercomputer software distributor.

One of the most interesting aspects of Japanese supercomputing usage is that there are four commercial sectors in which Japanese companies own supercomputers but where there are no supercomputers in analogous U.S. companies. These sectors are construction, steel, public utility, and transportation. In these latter two areas there are actually very few companies involved, but the supercomputing usage is unique and so we assigned them to distinct categories.

The transportation sector consists of only a single company. East Japan Railway (JR), the largest of six railway companies created when the Japanese National Railway was privatized, has a four-processor CRAY Y-MP vector system that it uses for crash simulation and other structural analyses. These applications are obviously quite similar to those of the automobile industry, but to the best of our knowledge, JR is the only railway company in the world to have its own supercomputer.

Perhaps the most novel usage of supercomputing relative to the United States comes from the Japanese construction industry. Construction is the largest industry in Japan, accounting for nearly 10% of the country's GNP, and it is about three to four times the size of either the automobile or the steel industries [54]. The Japanese construction companies participate in nearly every aspect of construction work, from design to actual building, and projects range from residences and offices to bridges, tunnels, airports, power plants, and off-shore structures. One Japanese construction company played an important role in building the Suez Canal (probably without using a supercomputer).

The construction companies that own supercomputers are among the largest in Japan: Kajima, Taisei, Penta-Ocean, Tokyu, and Ohbayashi. In 1993 these companies had annual sales ranging from about U.S. \$3 billion (Tokyu) to about U.S. \$20 billion (Kajima and Taisei). Although Japanese construction companies devote much less of their resources to R&D as a percentage of sales than do most other Japanese companies (about 5% as opposed to about twice that much for electronics and manufacturing firms), they still spent between U.S. \$38 – 250 million on R&D in 1993, which is reportedly more than U.S. construction companies spent [54].

Of the six supercomputers used by Japanese construction companies, all are vector machines and all but one, a single-processor Cray Research system at Tokyu Construction, are Japanese-manufactured. Two companies, Taisei and Ohbayashi [55], have very powerful vector systems, similar in performance to machines installed at nuclear research facilities in Japan or at aerospace companies in the United States.

In general, R&D in the construction industry focuses on environmental issues; materials; sound, wind, and vibration studies; and analysis of construction methods. Some of the types of simulations carried out are [54, 55]:

- structural analysis of buildings, for both active and passive support during seismic events;
- flow of heat and air in buildings, including prediction of temperature in air-conditioned rooms;
- diffusion of smoke from automobiles, stacks, and from fires;
- effect of high-speed vehicles entering tunnels;
- ocean waves, including those created from earthquakes;
- geophysical studies of foundations;
- acoustical analysis in amphitheaters.

In contrast to supercomputing usage in other commercial sectors, construction industry simulations are not focused as much on cost reduction or rapid product development. Rather, they are intended to provide data that would be difficult or impossible to obtain otherwise. An important example of this is fatigue evaluation of existing structures such as bridges. On-site testing is expensive, and more importantly, it must be done non-destructively, so that data at or near the actual stress limits cannot be obtained at all.

Construction industry simulations are therefore seen as a possible means of increasing structural safety, which is apparently an area that has been plaguing Japanese construction companies recently [56]. The problem that arises with construction industry supercomputing is that in the absence of extensive experimental data, the reliability of the simulations is often called into question. Even though some of the Japanese construction companies were among the first supercomputer users in Japan, beginning around 1986, much of their work is still focused on validation of their computer models. In our in-depth profile of one of the Japanese construction companies in the next chapter, we note that most engineers still refuse to accept the supercomputer results in the absence of some experimental data.

The Japanese construction industry research in high-performance computing produces results that are impressive in terms of their obvious practical application and also often in terms of their ambitious goals. As an example of what we mean by "ambitious goals," consider that Taisei is attempting to calculate the temperature comfort level of a person in an office as a function of such detailed parameters as the person's clothing and metabolic rate, in addition to building-design variables such as air flow and room geometry. Related work is carried out by automobile companies, which are interested in roughly the same kinds of information for car passengers, but from what we have seen, these simulations are much less detailed.

In the U.S., large construction companies are generally more restricted to "construction management" activities and there is less interest in basic research and computer simulation than there is in Japan. A case study of Bechtel given in Chapter 3, confirms this. Some of the structural analysis work required by large U.S. companies is sub-contracted to small firms whose size clearly precludes owning even a small supercomputer. We presume that if such companies were interested in supercomputer simulations they could collaborate with university [57, 58] or national laboratory [59] mechanical engineering departments. However, for the most part they are not doing so and are using workstations and PCs instead. Using these less-powerful machines implies enormous simplification of the computer model and therefore less accurate simulations.

It is difficult to discern the effect, if any, of supercomputing on the Japanese construction company corporate bottom line. These companies may have purchased their own supercomputers because industry-university collaborations are more difficult to implement in Japan. Also, in Japan there is a strong "follow-the-leader" tendency, wherein a company decides to implement a given technology simply because another has. In Japan it is possible that this attitude transcends industrial sector borders; that is, construction companies decided to implement supercomputing because so many other industries in Japan did in the late 1980s.

One observer in Europe believes that the design and construction industries are in a transition state between that of a low-technology, labor-intensive enterprise to one that is compute-intensive and driven by high-technology [60]. There is little doubt that the Japanese companies are well-positioned to incorporate such high-tech features into their business, probably more so than their American or European competitors. However, this same observer believes that the main technological improvements benefiting construction industries will be in non-supercomputer areas such as communications (networking), robotics, workstation-based CAD systems, and even areas as remote from supercomputing as hand-held and pen-based computers.

Nevertheless, we conclude that supercomputer simulations of physical structures is an area that should be more actively encouraged in the United States. However, because such simulations generally do not result in any financial gain, but rather improvements in safety, the benefits will accrue more to society as a whole than to an individual company, and the government is probably the only means of continuing support.

The steel industry is another one in which there appears to be more supercomputing activity in Japan than there is in the United States, at least based on the number of installations. Three companies, Sumitomo Metals, Kawasaki Steel, and Nippon Steel, have low-end vector supercomputers from Fujitsu or NEC. In the U.S., we know of no steel companies that have their own machines, although U.S. Steel has been a corporate affiliate of the Pittsburgh Supercomputer Center (PSC) since 1989, and its Technical Center researchers use the Cray Research systems installed there.

Traditionally, supercomputer simulations in the steel industry are used to model flows of the molten materials, analyze stability of solid structures, investigate chemical reactions, and explore material properties. Quantum chemistry simulations are even performed on steel products consisting of complex mixtures of various substances, and electromagnetic simulations aid in the analysis of molten steel couples.

However, in Japan, not only is supercomputing (and other advanced computational techniques such as expert systems, fuzzy logic, neural computing and robotics) used as an important tool to aid R&D related to primary steel products, but it is also becoming something of an end in itself. It has been noted that whereas steel-making in general is not an industry expected to experience significant long-term growth, computing or information science is such an area, and thus Japanese steel makers are attempting to diversify so as to pursue computer-related fields [61]. For example, a Japanese steel industry group called Japan Iron and Steel Foundation is one of the partners of Japan's Real World Computing Partnership [62]. Also, one Japanese company, Nippon Steel Tubing Company (NKK), has an agreement with Convex Computer in which the former will send a dozen or so software developers to the computer company's headquarters in Texas to learn about their new parallel processing system, with the hope of eventually developing third-party software for it [61]. And Sumitomo Metal Industries, a leading integrated steel maker, has been attempting to enter businesses such as electronics and biomedical science, as well as computer-related areas such as systems engineering and measurement and control [63]. Sumitomo has begun selling integrated circuits and it uses its NEC SX-2 supercomputer to carry out device simulations.

As noted above, the computational resources available to researchers at Japanese steelmakers are rather limited in terms of absolute performance, consisting of vintage 1986-1990 single-processor vector machines. In contrast, researchers at U.S. Steel have access to some of the largest and most recent Cray Research systems at PSC. This would include one of the largest massively-parallel Cray systems installed anywhere, although because it is so new, it is not known if U.S. Steel is using it.

One observer in Japan has suggested that not everyone within the Japanese steel research community feel they need additional computational power [64]. A foreign researcher in the Metallurgy Department at the University of Tokyo has stated that overnight turn-arounds for simulations are sufficient and that a more important goal is wider availability of systems and simplicity. Most likely, the U.S. Steel researchers are using single-processors of the PSC CRAY C90 system, which does not give them any faster turn-around on their jobs than their Japanese counterparts. U.S. steel researchers probably do relish the fast turn-around time the Cray gives them relative to their in-house computing systems (workstations) but it is not clear that the vastly superior resources afforded by PSC give U.S. Steel a distinct advantage over the Japanese competitors. It is quite likely that the computer models, which include turbulence to model the flow patterns of the molten steel, are the limiting factor, not computer power.

Finally, there are two supercomputers in Japan located in what we have called the "utility" sector. The two sites are Tokyo Gas and the Central Research Institute of the Electric Power Industry (CRIEPI), a non-

profit private corporation whose support comes largely from several electric power companies and the Japanese Government. CRIEPI has a small Thinking Machines, Inc. parallel computer which is used for global-scale numerical simulations such as air quality studies. The application of interest to Tokyo Gas is not known. The machine there is a single-processor Cray vector computer with very large main memory, so most likely, some sort of three-dimensional computation is being performed. No power companies in the U.S. are known to have supercomputers of their own, although, again, university collaborations probably exist. Other power companies that do have supercomputers are Hydro-Quebec and the French company Electricite' de France [65]. The latter is the world's largest electrical utility, and it owns two Cray Research vector machines and a small Cray Research parallel system, i.e., a significant amount of computing power. Numerical simulations of interest to electric power companies in general include structural analysis, electromagnetics, power distribution analysis, and climatology. Additionally, Electricite' de France is involved with nuclear physics and nuclear power plant design.

2.4.3.4 Effect of the 1994 Japanese Government Procurements

Now that we have described the current supercomputing usage in Japan it is important to explain some of the circumstances of its evolution. Specifically, during the Japanese fiscal year 1993 (which ended in March 1994), three Japanese Government agencies purchased a total of 16 supercomputers to be placed in universities and various government laboratories [66]. Some of these machines were part of the normal budget process but 11 were purchased under a special "supplementary" budget that was appropriated as part of an economic stimulus package. *Our data show that these 16 machines account for 35% of current total supercomputing power in Japan.* In addition, by the end of the 1995 fiscal year, these agencies anticipate the purchase of eleven more machines [67]. Table 2.7 lists the machines, their performance, and where they were installed (an asterisk denotes some of the machines expected for FY95).

Because supercomputer peak speeds are increasing dramatically as a result of massive parallelism, it is not really so surprising that much of Japan's supercomputing power comes from installations of recent parallel machines. More important, however, is the contrast with recent installations by Japanese companies during Japan's FY93, an approximate guess at which would be seven or eight machines, with another five being installed in-house at the vendors. This estimate is probably low. In any event, the combined power of these new commercial-sector machines is about 28 GFLOPS, even counting the in-house machines, but this would compare with nearly 400 GFLOPS installed by the government during its 1993 fiscal year. Thus, the Japanese Government is now the primary funding source for supercomputer purchases in Japan, a situation that is quite different from that of a few years ago.

It is important to examine why the Japanese government decided to earmark supercomputers as part of its economic stimulus program. On the face of it, there was simply the desire to increase the available computing power. It has been reported that even in early FY93 Japanese governmental agencies wanted to buy supercomputers in order to double the number of machines at national research laboratories, and because it saw the machines as a "social asset" [68].

However, there are probably other reasons. One may be an attempt to improve the quality of basic research in Japan. Japan as a whole has been criticized for relying on fundamental science advances made by basic researchers elsewhere and financing product development instead [69]. Now, however, Japanese funding for basic research seems to be growing substantially [69], and supercomputing is seen as an important component. Traditionally, basic research in Japan is carried out largely within companies, but there is a perception in Japan that this gives the U.S. an advantage because in the U.S. such research is generally done at universities. Japanese industry R&D leaders have called upon the government to strengthen academic research, and perhaps part of the response has been to place some large supercomputers at universities [70].

Funding large supercomputing centers as part of the supplementary budget procurements is also seen as an opportunity to counter both the technical and public relations gains made in the U.S. by the High Performance Computing and Communications (HPCC) program. Recently, Japan's Science and Technology (STA) agency announced the start of a FY95 "computational science project" which encompasses some of the supercomputer procurements [71]. The description of the project sounds very much like an attempt at a formal (albeit very small-scale) version of the U.S. HPCC program. Another successful model for Japanese emulation is the U.S. National Science Foundation Supercomputer Centers.

The possibility of easing trade friction with the U.S. by choosing U.S.-manufactured systems for some of the sites may have been another contributing factor. The procurements were under intense scrutiny by U.S. Government officials to ensure that American-made machines were considered fairly [72]. In fact, six of the 16 machines chosen were from U.S. vendors (Thinking Machines, Intel, MasPar, Digital Equipment, and Cray Research). Recall that in the beginning of this chapter we suggested that competition with Cray Research comes at least as much from other U.S. parallel manufacturers as it does from Japanese manufacturers. The Japanese supplementary budget supercomputer procurements are, therefore, a good example of this. One puzzling aspect of the procurements is the selection of an Intel Paragon by the National Aerospace Laboratory (which already owns the NWT). There had been some bad press in Japan regarding the Intel Paragon installed at the Real World Computing Program [73]. The Paragon implementation of parallel processing is different from that of the NWT, and so perhaps the Paragon was purchased to give NAL researchers experience with other machines and other programming methodologies. In terms of capability or capacity, however, NAL's Paragon is probably considerably less useful than the NWT, and is therefore somewhat superfluous as a computing engine.

The Japanese Government has played an important role in the development of the Japanese supercomputer industry by helping to form large research collaborations (such as "Superspeed," Fifth Generation, and Real World Computing; see last year's report [8]). As an economic stimulus, the supplementary procurements enlarge the role of providing income for Japanese supercomputer manufacturers. The question that arises is whether this increased role is permanent or temporary, i.e., whether an improvement in economic conditions will mean a decreased government role and increase in civilian supercomputer purchases similar to the level of a few years ago. The "downsizing" of the supercomputer industry, in which less expensive workstations are favored over large-scale machines,

suggests that increased government support will become more important.

Another important aspect of these latest procurements is that they were the first public procurements involving highly-parallel systems. Ten of the winners so far are parallel processors, of which Japanese-manufactured ones account for five. There have been suggestions that the Japanese Government uses its procurement powers in order to nurture new Japanese industries [74]. Although Fujitsu is obviously a well-established company, its new line of massively-parallel computers could be viewed as being a nascent "industry" in need of nurturing; the supplementary budget may have been a way of doing so. Note that the money spent on the development of the NWT is effectively an additional guarantee to Fujitsu.

Table 2.7: List of Supercomputers Installed or Expected to be Installed as Part of Japanese Government Purchases During FY93 and FY94. Data from References 66 and 67.

Machine	Site	Type	LINPACK Rating (GFLOPS)
Fujitsu VPP500/30	Angstrom Technology Partnership	P	33.0
Fujitsu VPP500/28	Institute of Phys. and Chem. Res. (RIKEN)	P	30.8
Fujitsu VPP500/10	Communication Research Laboratory	P	11.5
Fujitsu VPP500/7	Institute for Space and Astronaut. Science	P	8.3
NEC SX-3/44R	Tohoku University	V	23.2
NEC SX-3/34R	Institute for Molecular Science	V	17.4
Hitachi S3800/3	Institute for Metals Research	V	5.2
Hitachi S3800/1	Meteorological Research Institute	V	1.7
CRAY C90/16	AIST Research Information Processing System	V	13.7
TMC CM5E/128	Angstrom Technology Partnership	P	12.8
CRAY T3D/128	Power Reactor and Nuclear Fuel Development Corp.	P	10.7
CRAY Y-MP/M92	National Aerospace Laboratory	V	0.55
Intel Paragon	National Aerospace Laboratory	P	9.8
MasPar 2216/16k	National Cancer Institute	P	1.6
MasPar 2216/16k	National Cancer Institute	P	1.6
Fujitsu VPP500/16	Tokyo Univ. Institute for Solid State Physics	P*	18.4
Fujitsu VP-260E and Fujitsu VPP500/16	Kyoto University Computer Center	V* P*	5.0 18.4
CRAY C90/16	Tohoku U. Research Institute for Fluid Science	V*	13.7
Cray C916/12	Tokyo Institute of Technology	V*	10.3
Hitachi S3800/3	Hokkaido University	V*	5.2

P = Parallel; V = Vector

3. Presentation of Case Studies

3.1 Introduction

In order to determine first-hand the role that computational science plays in Japanese research and development we undertook field research at several Japanese companies that have their own supercomputers. Our intention was to survey a variety of factors related to supercomputing usage such as those listed below.

- history of computation at the company.
- allocation of computational resources among company researchers
- the extent to which simulations are integrated with other techniques
- what areas of science/engineering are investigated
- which computer programs are used and what is their origin
- what kind of networking facilities are used with the supercomputer
- what level of performance is achieved on the supercomputer

We also hoped to learn about specific examples in which the company had applied supercomputing technology to a problem and had directly obtained a cost savings or reduction in design time as a result. It was interesting in itself to know if companies kept track of such instances and indeed, to what extent they kept track of any of the areas about which we queried them. We got the impression strongly at one site, and sometimes at others, that the information they showed us was collected specifically at our request and in fact, we may have done the company a service by causing them to examine their system as we wanted to.

One of the important decisions we had to make early in this study related to coverage, i.e., whether we would attempt a comprehensive survey across all major sectors in Japan or concentrate on a few key sectors instead. Another variable we considered was the type of supercomputer used at the facility. We have a great deal of familiarity with Cray Research supercomputers, but much less with the Japanese supercomputers, especially those from Hitachi, about which we know very little because their machines are not sold in the United States.

Ultimately, the entire decision was made moot because obtaining invitations to visit companies was more difficult than we had expected. In effect, the sectors and machines we covered were determined by which companies we were able to contact and which of those agreed to host us.

The companies we visited are listed in Table 3.1. Several important sectors are missing, such as chemical, pharmaceutical, financial, and "heavy industry" companies such as Mitsubishi or Kawasaki. Also, the only company with a Hitachi supercomputer is Hitachi itself.

Case Studies of U.S. Supercomputing Installations

To better understand the correlation between supercomputing in Japanese industry and comparable companies in the United States, we have investigated, where possible, similar American installations. Because of time constraints and financial concerns, most of the study for these companies has been conducted through telephone interviews or through correspondence. Because of this, it is not as thorough as the personal visits conducted in Japan.

Table 3.1. Brief Profile of the Companies Visited.

Company	Date	Industry	Type of Supercomputer(s)
Matsushita	5/9/94	Electronics / Semiconductors	Cray / Fujitsu
CTI	5/10/94	Computer Services	Fujitsu
Toyota Central Research and Development Lab., Inc.	5/11/94	Automobile	NEC
Toyota Motor	5/12/94	Automobile	Cray / Fujitsu
Nissan Motor	5/13/94	Automobile	Cray
CRC	5/16/94	Computer Services	Cray
Taisei	5/17/94	Construction	Fujitsu
	5/19/94		
Hitachi	5/18/94	Computer Vendor / Semiconductors	Hitachi
JRCAT/RIPS	5/20/94	Non-Defense Gov't Research	Cray / Fujitsu / TMC

3.2 Matsushita Electric Industrial Company

Matsushita Electric Industrial Company, LTD, is one of the world's largest consumer electronics manufacturers, and is the world's leader in VCR production. The company and its subsidiaries (314 companies) had about U.S. \$60 billion in revenue in 1993 and about 250,000 employees worldwide. Matsushita spent about U.S. \$4 billion (5.7% of sales) on research and development in 1992, which is the second largest amount of any company in Japan (Toyota is the largest).

Our visit to the company on May 9th consisted of visits to two separate laboratories, the Central Research Laboratory and the Semiconductor Research Center. The Central Research Laboratory is located midway between Osaka and Kyoto in the new Kansai Region Science City. This is a new collection of research centers in a rural area called Keihanna, and other occupants include various government laboratories such as the Research Institute for Innovative Technology for the Earth (RITE) and the Advanced Telecommunications Research Institute (ATR). The Kansai Research City, as it is also known, is an area in which Japan's Ministry of Posts and Telecommunications intends to implement a pilot project soon to study next-generation telecommunications technologies. Approximately 137 companies, probably including Matsushita, will participate in fiber-based broad-band ISDN experiments.

The Matsushita Central Lab is in a new, very modern-looking, and quite spacious building, in which about 300 people carry out both experimental and computational research in areas such as materials, advanced science, intelligent electronics, ultra-precision machining, thin films, and health electronics. There is also a "Lighting Research Laboratory."

Matsushita has a four-processor air-cooled CRAY Y-MP supercomputer, which is used for two basic kinds of simulations in roughly equal proportions: electromagnetic (EM) field analysis (solution of Maxwell's equations using a Finite-Element approach) and quantum chemistry for materials research. All of this work is done by about ten researchers, all of whom can access the Cray machine from their homes, using a PC and modem to dial into the Matsushita network. As we learned later in the week, this is a rarity in Japanese companies.

In order to find out about the usage of the Y-MP we talked to Dr. Shin-Ichiro Hatta, whose title is listed simply as "Manager, Central Research Laboratories." Hatta showed us the results of two simulations related to EM, a microwave oven and a rice cooker. Matsushita uses simulations to optimize the design of these products, because by simulating the emissions of the EM sources, the temperature distribution in the microwave oven chamber or in the rice maker can be calculated. Using the calculated temperature distributions the scientists can adjust the design of the unit to accomplish various goals, such as uniform heating. The results of the simulations are also compared with Matsushita's careful experimental measurements. Hatta claimed that the Y-MP simulations are both less expensive and more accurate than the experiments. The rice makers use a new kind of heating device altogether. Previously, they used simple resistance heating but the new versions, which sell for upwards of \$300 each, use induction

heating systems, in which oscillating electric fields induce Eddy currents and thus heating in the metal-lined rice container. Hatta stressed that the principal value of the supercomputer in these cases is faster design time for the products. Interestingly, an independent source mentions that defective Matsushita rice cookers are becoming a significant problem, making enhanced quality control an immediate issue for the company [2].

From other reports we have seen, we know that Matsushita is doing research related to neural networks and their use in factory automation. The glossy literature describing the Central Research Laboratory we obtained during our visit states that neural-fuzzy controllers are also used in the new rice makers, so vector supercomputers are not the only modern computing technology that Matsushita is using to improve its consumer products.

Dr. Hatta also spoke about the second principal area in which the Cray is used, which is in various kinds of quantum chemistry simulations related to potential high-temperature superconducting materials such as doped CaCuO_2 . Clearly, this research is more fundamental or "pure science" in nature, in contrast with the highly product-oriented EM simulations discussed above. Hatta has published some of his results in this area with collaborators from DuPont. The codes used for this work implement the density-functional method. They generally vectorize very well, and thus their performance on Japanese-manufactured supercomputers, such as the NEC SX-3 or Fujitsu VP2400 might be considerably better than on the CRAY Y-MP that Matsushita owns.

All of the computer programs used in both the product simulations and the materials chemistry were developed "in-house" by Matsushita scientists. Dr. Hatta, who is a physicist, stated that little time is spent by the scientists optimizing the codes, beyond that which can be obtained quickly and easily. Hatta also believes that further advances in simulation techniques are more important than faster computers, a sentiment, incidentally, that was echoed by another Matsushita scientist with whom we spoke later in our visit. Nevertheless, Hatta stated that he would like to persuade his managers that more processing power is needed at Matsushita. Currently, none of the Matsushita codes utilize more than one processor (at a time) of the Y-MP, and runs are usually carried out overnight.

At the Semiconductor Research Center we learned that research in semiconductor technology is a central part of Matsushita's R&D effort for two reasons. First, integrated circuits are now at the heart of many Matsushita products, and the number of applications is probably growing. Second, a broadly-defined category called "electronic components," of which ICs are a principal part, now accounts for about 12% of Matsushita's sales.

Matsushita's Semiconductor Research Center (SRC) is located in Osaka, right across the street from the company's headquarters. The lawn in between the two has a fountain and several statues of famous scientists/inventors who contributed in some way to the electronic device world. The central and largest statue is of Thomas Edison. The buildings in this area, including the SRC, are late 1950s - early 1960s vintage, and show a fair amount of age and wear.

Our principal guide for a tour of the SRC was Dr. Shinji Odanaka, a very bright and intense young scientist who is with the Matsushita VLSI Technology Research Laboratory.

The Matsushita Semiconductor Research Center has two high-performance computing resources, both of which are used primarily for device simulations (similar to those performed by the well-known SPICE program, but in-house codes are used at Matsushita). The first of these is a Fujitsu VPX240, which is a 1990-vintage single-processor multi-pipeline vector supercomputer with about 2.5 GFLOPS peak performance. The VPX240 is only the middle-range version of the Fujitsu VPX series, which is a subject we discuss in the section below. Matsushita's VPX240 was purchased as an upgrade to an earlier Fujitsu vector supercomputer owned by Matsushita, the VP-200. The "VPX" designation means that the machine is a Fujitsu Unix (UPX) version of Fujitsu's VP2400 machine.

The second supercomputer used at SRC is one that was developed by Matsushita itself. It is called ADENART, and it was developed by Matsushita in collaboration with the well-known Professor Nogi of Kyoto University. Several years ago Matsushita began exploring the possibility of entering the computer business and ADENART was to be its initial product. ADENART is one of a small number of massively parallel computing projects underway in Japan. The company has apparently decided, however, that building and marketing the ADENART is not an option it will pursue any time soon. Nevertheless, one prototype was built and it is now used to carry out Monte Carlo particle transport simulations. It also runs a parallel version of the same vector device simulation code that is run on the VPX240. Dr. Hiroshi Kadota, Matsushita's principal scientist on the ADENART project, told us that Monte Carlo performance of the ADENART is about the same as it is on the VPX240, and since it is highly vectorized, we would guess that sustained performance is a little less than 1 GFLOP.

Another important use of supercomputers at Matsushita was revealed to us during discussions following our tour of the SRC. Apparently either the VPX240 or the CRAY Y-MP (it was not clear which) was instrumental in the development of Matsushita's flat-panel television screen. As far as we were able to tell, the flat-panel screen uses a large number of electron beams scanning a much smaller but much more precise area than does a normal TV screen and the supercomputer simulation is required to align the beams.

Matsushita is using a supercomputer in a way that directly affects its competitiveness in consumer products. A multi-million-dollar machine is being used to optimize the design of a product that will sell for at most \$300 or so. On the other hand, the same supercomputer is being used for simulations in basic science, in an area that probably has no specific product payoff for at least ten years, yielding results that are published. Also, supercomputer simulations are being used to help design a product of absolutely enormous commercial value, namely flat-screen video. This is an area in which the United States has virtually no market presence currently, but the (U.S.) government is considering taking an active role in industrial development of a flat-screen product.

All of these simulations are carried out on a Cray supercomputer using software developed in-house by Matsushita. This makes us wonder why Matsushita bought a Cray machine to begin with, since one of the most oft-quoted advantages of Cray supercomputers world-wide is the large volume of applications software that is available for them. It is a question we did not ask our hosts explicitly, since we believed the issue of why a given entity buys a given machine to be too politically sensitive.

We believe there is evidence that Matsushita's progress in the application of high-performance computing to its R & D projects is being hampered by economic considerations. In the semiconductor research area, Dr. Odanaka stated that even with two supercomputers, he does not have enough power to run the simulations he would like as fast as he would like. He stated that something closer to about 10 GFLOPS performance would be needed. In view of Dr. Odanaka's statement about the overwhelming importance of device simulation at Matsushita, we found this admission surprising. There are several machines available that would deliver the required performance, including the more powerful version of the Fujitsu vector supercomputer they already own. Our conclusion is that the company is probably unwilling to spend the money required to purchase the additional computing power. The effects of the Japanese recession and strong Japanese currency worldwide have taken a significant toll on Matsushita profits in the last few years, and this could be behind the reluctance to purchase a more powerful machine.

As a final thought, we asked a collection of Matsushita scientists what technology improvement could most significantly benefit the company's computing-based R & D in the future, and Dr. Odanaka stated emphatically that it would be improvements in networking technology. Apparently, the design of semiconductor ICs is a multi-stage process that can make use of existing data regarding solid-state characteristics at several points along the way. The ability to share these huge databases easily and quickly between researchers is a break-through that Odanaka looks forward to eagerly.

3.3 Computer Technology Integrator Co., Ltd. (CTI)

CTI is a service company whose biggest recent contract was for design of components of the Boeing 777, in cooperation with the Seattle-based company. This project is now complete, and the company is seeking other large projects to replace it and to augment their existing work with several Japanese power companies. The company expects to get a portion of the development work for Japan's new fighter aircraft.

Five Japanese subcontractors, including Mitsubishi Heavy Industries, Kawasaki Heavy Industries and Fuji Heavy Industries, have been involved in the huge Boeing 777 project. Since the 1970's, Japan has been an influential partner in the aircraft design and production business. In the production of the Boeing 767, Japanese investment in the project accounted for 15% of the value of the fuselage; this is expected to rise to 21% for the 777. Even more significant is expected to be the effort put into the proposed super-jumbo jet project, an effort to design and build an airplane specifically suited for the needs of Japanese cities, airports, and travelers. No doubt CTI is hoping to provide a large part of this project's computational requirements.

CTI employs approximately 300 persons. Three years ago the company moved into its present facility, a multi-story, modern building located in the outskirts of Nagoya. The building is very impressive, and has been built subject to the MITI standards of earthquake protection, floods, and other natural disasters. One strange feature of the building, however, is that it contains almost no windows. The company is young and appears to be enthusiastic in its outlook. The major technical arm of the company, the Science and Technology Division, is divided into 3 sections:

- Computational Science --- includes CFD, environmental studies for government agencies (including biosphere); effects of various development plans on estuary and coastal regions; material science; and structural analysis.
- Electric System Power Analysis --- provides operational and computational application support to power systems built by Chubu Electric, a major investor in CTI.
- Multimedia --- includes TV animations, and scientific visualization and CAD development.

The major computational resources at CTI are a Fujitsu VP2400 supercomputer, front-ended by an IBM 9021 mainframe, and numerous workstations from vendors such as Apollo, Sun, SGI, Fujitsu, and IBM. For the VP2400, CTI runs MSP/ex, FORTRAN77 and Fujitsu math libraries and the PHIGS graphics library. MSP/ex is Fujitsu's proprietary operating system, which has the "look and feel" of IBM's MVS system, and is considered archaic and very difficult to use by most U.S. scientists. Their programming tools, VECTUNE and FORTUNE, are provided by Fujitsu. The choice of a VP was based primarily on the fact that CTI's biggest customers were using the Fujitsu M400 and their codes would run unchanged on the VP. Thus, CTI provides an important example of a company that purchased a Japanese-manufactured supercomputer in order to maintain compatibility with existing mainframes.

The most heavily used code is NASTRAN, which is used for structural analysis. CTI's version of NASTRAN comes from McNeil Schwindler, who is the main supplier in the U.S., but we were told that typically Fujitsu customers receive less support for NASTRAN than Cray users do. Other codes used are ABAQUS (nonlinear structural analysis); FLUENT (thermodynamic analysis); STAR-CD from Britain (also used for thermodynamics); STREAM, an easy-to-use code developed in Japan for thermofluid dynamics analysis that does not deal smoothly with complex geometries; USAERO, used for the simulation of a train entering a tunnel (unsteady flow analysis-BEM); COSMOS, an in-house code for environmental analysis; and MASPHYC/MD, developed by Fujitsu for molecular dynamics.

The VP2400 is configured with two scalar units (SUs) and one vector unit (VU), an unusual configuration for a supercomputer. The way such a machine works is that the scalar units "feed" the vector units when a vector processing job is run, and they operate on their own when only scalar jobs are run. Optimal performance happens when both scalar units and the vector units are kept busy all the time, and this would be indicated by 200% utilization. Statistics collected by CTI indicate from 60% to 160% machine utilization (including the VU and both SUs). Thus, the extra scalar unit appears to be paying off for CTI.

Jobs are submitted to the VP supercomputer through the front-end machine, the IBM 9021-500, which shares disks with the VP. Disk capacity is 70 GB, of which 45 GB is shared with the IBM. CTI also generates income by selling time on the VP, and primary users include Matsushita (Panasonic), Fujita, and Mitsubishi Heavy Industries, which has the largest share by far, with 245 users. Generally, jobs have a short execution time. The average job duration is 3.4 CPU minutes; maximum job duration recorded in 1993 was 28+ CPU hours. The VP can be used by remote access (such as Seattle). Data from Seattle can be submitted over encrypted lines; however, use from an employee's home, as is done at Matsushita, is not allowed.

CTI considers its graphics expertise to be its strength. The company is currently engaged in producing television commercials and graphics for use with local TV shows. They are also discovering that many customers are more interested in using powerful workstations rather than a supercomputer. This will no doubt impact their future business, as these same customers may be purchasing their own workstations and utilizing CTI in a different, and possibly, less profitable, manner. They see possibilities for future work in the multimedia area and are pursuing this strongly; it appears that this is the company's main hope for continued existence.

3.4 Toyota Central Research and Development Laboratory, Inc.

Toyota Central Research and Development Laboratories, Inc. (TCRDL) was established in 1960 to do basic research for the nine companies in the Toyota Group. It is located in several 1950s-modern buildings in Aichi Prefecture near the city of Nagoya. Research fields include automobile-related projects, advanced materials, material analyses, computer technology, communications, and environmental issues. As a service and research company, money to support projects at TCRDL comes from the companies in the Toyota Group, and at least one-half of the research is directly connected to the needs of the Toyota Motor Corp. (TMC).

Dr. Tsuguo Kondoh, who is a Senior Researcher and Manager of TCRDL's Applied Mathematics & Physics Lab., told us that prior to 1991, computational science at TCRDL was done on a NEC AC0900 mainframe and on leased time on supercomputers at two external organizations: Japan's Institute for Computational Fluid Dynamics and Recruit, Ltd. Management was finally convinced to buy a supercomputer for computational science use at TCRDL, but it took five years to convince them that it was a good idea. In 1991, they purchased an NEC SX-3/14T. An important justification in the decision to purchase the SX-3 was the decrease in time taken for simulations to run. The designation of the model as an SX-3/14 means that it has a single processor and four sets of arithmetic vector pipes. We were told that the "T" designates a special model ("T" for Toyota) built especially for Toyota that has extra hardware for gather-scatter operations.

TCRDL is primarily engaged in two areas of computational research for Toyota, computational fluid dynamics (CFD) and electronic structure analysis. The code used for CFD, called FIRE3D (Flow In Reconfigurable Engine), was developed at the TCRDL for the SX-3. FIRE3D is now used as a production code by Toyota Motor Corp., and is also thought to be useful to other companies outside of the Toyota Group. According to the researchers at TCRDL, it may be marketed as a commercial product. However, we believe it is the only code developed by TCRDL now used in production by the parent company. Currently the performance of the code is a very impressive 3 GFLOPS on the NEC SX-3/14. Last year, FIRE3D simulations ran for a total of nearly 4500 hours, which is about 50% of the available machine time.

Eighty percent of the time used on TCRDL's SX-3 is given to CFD simulations, with the remaining 20% of the time used for electronic structure calculations in new materials research. Even with the very fast SX-3, it still takes nearly six months to get CFD results for a complete study of, for instance, drag coefficients on various body shapes. This represents only about 100 hours of time on the SX-3, so the real time-consuming part of these studies comes from the human-computer interaction, in other words, grid generation. Other studies using a CFD approach are the simulation of air flow over the sunroof of an automobile and air conditioning flow inside of an automobile.

The scientists at TCRDL cited several examples of computer simulations helping them gain knowledge that would not be available by other means. Two of the examples are from CFD and one is from electronic structure calculations. These research simulations involving computational fluid dynamics are used to drive research directions and are not yet used for design decisions on the automobiles. Examples:

- Details of fluid motion around the body of an automobile that cannot be determined by wind tunnel tests;
- Fluid motion in a torque converter. A torque converter is the mechanism for an automatic transmission and, because the automatic transmission fluid is sealed completely inside an iron cover, observation of fluid motion from outside is not possible;
- Electron density distribution in semiconductors. We were told that this simulation is useful because, experimentally, data are available only for silicon crystals due to the high quality of crystal needed for such measurements.

Although the simulations are currently done on a single processor of an SX-3, economic and time constraints are forcing the scientists to consider parallel processing. Kondoh thinks that multiple processing on an expanded SX-3 system would be one path to take, because the codes now run on a single processor and decomposing for a shared-memory system would require the fewest number of changes. However, TCRDL has also looked at distributed-memory systems for one of the codes. A first-principle electronic structure code written by Dr. Hidemitsu Hayashi, who is also a researcher in the Mathematical Physics Group has been run on a cluster of ten engineering workstations. Curiously, rather than use a publicly available communications package such as PVM (Parallel Virtual Machine, a very commonly used system for this kind of development in the U.S.), Dr. Hayashi developed an inter-processor communications system himself. The code on this system runs ten times slower than it would on the SX-3, but it costs ten times less. When we asked why they developed their own communications package instead of using a publicly available one such as PVM, the researchers said that there was fear that the U.S. Government would restrict export of such software to Japan. This seems to be a driving force behind other in-house development of software as well.

Advanced materials research at TCRDL is motivated by the need to reduce vehicle weight and improve engine performance in Toyota vehicles. One area of research area that has won TCRDL scientists several publication prizes is Tribology, which is the integrated analysis of machines, materials and lubricating oil in order to effect energy savings and reliability. In addition to these two fields, scientists are researching semiconductor and advanced information techniques, microelectronics (using silicon semiconductors), optoelectronics, human engineering, environmental chemistry, and biotechnology. These last research areas use primarily engineering workstations for computations, but the use of supercomputers is gaining acceptance.

During our tour of TCRDL computer facilities, we noticed that the machines were in a huge computer room that was nearly empty. Whether this is the result of long-range planning (build the room big now for future purchases) or the result of a miscalculation of how much high-performance computing would

be useful to this type of research is not known. Again, the economic downturn seems to have had an effect on this laboratory's ability to purchase supercomputers. There are many engineering workstations made by both American and Japanese manufacturers for the scientists' use, but they are all located in a special room. Scientists who need to use one have to make a special trip there to compute. This, it seems to us, would have a negative effect on computational productivity.

The scientists at Toyota Central Research and Development Laboratory consider their primary job to be basic research. They think that parametric studies, such as those using the NASTRAN structural analysis code, should be done by the companies such as Toyota Motor Company. The use of CFD has not been incorporated into automobile design at Toyota Motor Company, but they hope it will be in the future. These scientists also say that corporate decisions are not affected by the results of their simulations yet, but they hope, with time, to change this.

3.5 Toyota Motor Corporation

Following our visit to Toyota Central Research and Development Laboratory (TCRDL), we visited Toyota Motor Corp. (TMC) located in Toyota City, Aichi Prefecture. In addition to learning how simulations and supercomputers influence the operations at TMC, we also hoped to see how the company put some of the ideas we had seen at TCRDL into action.

Toyota Motor Corp., the leading member in the nine-company Toyota group, has a total of about 71,000 employees. Toyota City houses a multitude of Toyota automotive plants, encompassing design, testing, and production. Using 1992 statistics, TMC is currently the leader in car production in Japan and the world, surpassing General Motors by about 100,000 vehicles. Sales have gone down slightly over the past two years, due to a number of factors not discussed here.

A great many kinds of computational tools play a significant role in the design and styling of new automobiles, and supercomputing is just one of them. CAD systems also play a critical role. About 1,000 personnel at Toyota Motor use computers in some aspect of their scientific work; 300 of these are system analysts. In addition to CAD/CAE systems on smaller computers and workstations, the following supercomputer systems are used for various kinds of simulations:

CRAY X-MP - crash simulation

CRAY Y-MP - crash simulation, sheet metal formation

Fujitsu VP 211 - casting, CFD (computational fluid dynamics)

Fujitsu VP2200 - CFD, linear analysis

The major supercomputer software applications supporting these are:

Structural Analysis: NASTRAN; ABAQUS; DYNA3D; SURFES (local product)

Crash worthiness: PAMCRASH; CRASH (local product)

Fluid Dynamics Analysis: SCRYU; STREAM 2D; STREAM 3D (in-house codes)

EM Analysis

Dynamic Analysis: ADAMS

Acoustic Analysis: BEM

We were given representative timings for some of the simulations being done. Using an in-house engine gas flow code, one simulation takes 10 minutes on the VP2200 and 3 hours on an HP735. For crash simulation it takes 20 hours and 16 MW on the X-MP to simulate 80 milliseconds of crash; this takes 8 hours on a Y-MP. An actual physical crash simulation would cost approximately \$500,000; to develop a new car model, 45-50 simulations are usually done, so it can be seen that the use of a supercomputer in this area represents a true cost savings. The X-MP at Toyota Motor Company is 90% utilized, solving mostly engineering problems. When we asked if they had an interest in parallel computing, we were told

that some preliminary work in parallel computing has been done on the nCUBE computer. The company management, as one would expect, requires some justification for the use and purchase of supercomputers from a cost savings point of view. Because of the state of the economy at the present time, there are no current plans to upgrade the Cray machines to the newer, faster C90 model. It was quite clear that supercomputers provide a major resource in the operation and planning activities of Toyota Motor and that they are used skillfully and knowledgeably throughout the entire automotive design and manufacturing cycle.

While crash simulation is in production mode at the plant, CFD is still considered a research area. The most difficult aspect of CFD is, of course, the generation of the grid. For example, to simulate an engine compartment for use with CFD requires about six months of grid generation calculations. To simulate airflow around the entire car requires about three months to generate the grid. This seems to be an area that the company has great interest in and may determine supercomputer needs and acquisitions in the future. The most important issues regarding the use of supercomputing resources at Toyota are:

- mesh generation time

- cost and time reduction for new car models

- cost/performance

In spite of their status in automobile production in the world, we were told more than once that they do not feel that they are as advanced in the use of supercomputers as the Detroit companies. The economy has also hindered their efforts to investigate such innovations as parallel computing and a better understanding of supercomputers to help them.

3.6 Nissan Motor Corporation

Nissan Motor Corporation is Japan's second largest and the world's fourth largest automobile maker. Although the company had about U.S. \$60 billion in sales in 1993 there was a net operating loss for the year of about U.S. \$71 million.

Nissan carries out basic and applied research at its Nissan Research Center, located in Yokosuka, about 1.5 hours outside of Tokyo. We visited there at the invitation of Dr. Ryutaro Himeno, who is a Senior Researcher in computational fluid dynamics (CFD) at Nissan's Vehicle Research Laboratory.

Dr. Himeno joined Nissan in 1978 after receiving his M.A. in electrical engineering. He spent 1984 through 1986 serving as a researcher at Japan's Institute for Space and Astronautical Science in order to learn computational fluid dynamics (CFD), and for the work he did there, was awarded the Doctorate from the University of Tokyo in 1986. Himeno is well known both inside and outside of Japan. He is a frequent speaker at meetings abroad, and was the recipient of the Cray Research "GIGAFLOP" Award for sustained high performance on a Cray system. His work in CFD at Nissan has been published in a variety of places. A recommendation to speak with Himeno originally came from Dr. Myron Ginsberg, of General Motors Research Laboratory, who is well acquainted with Himeno and his work.

Nissan's high-performance computing resources are fairly varied and numerous. They have two Cray Research systems, a six-processor Y-MP with 64 MW of memory and a four-processor X-MP system. The Y-MP is used for the company's structural analysis and for Himeno's CFD work, and the X-MP is dedicated to crash analysis. There is also a four-processor Convex C-240, by now an old machine, which is used for the development of CFD codes.

Perhaps the most important aspect of Himeno's work is that he realized fairly early in his studies that application of CFD to Nissan's car design process is severely limited by CPU power. For that reason he adapted his code to run on a variety of parallel machines. He can easily use multiple processors of the CRAY Y-MP, at least when the rest of the company allows him to do so. However, Himeno is also actively searching for possible replacements for the Cray. He has published papers providing results on several parallel machines: an Intel iPSC/860 with 32 processors, an nCUBE2 system with 512 processors, and even the Matsushita ADENART system. In fact, Himeno told us that of available parallel machines today, he leans toward nCUBE if he were to buy a parallel system for installation at the Research Center. However, he also added, rather wistfully, that economic considerations will probably prevent such a purchase soon.

Himeno is also examining the possibility of moving his computations from the Cray Research machines to a cluster of RISC-based workstations. Nissan has experience running their codes on single-processor scientific workstations, such as those made by Hewlett Packard (HP). If a workstation cluster is chosen for the work, it would probably be one such as the IBM SP-2, Convex "Exemplar," DEC 7000, or an HP

system. Himeno has developed a benchmark test based on his CFD code using a fairly large grid (about one million points). He showed us benchmark results from several Cray Research and Convex systems. Interestingly, the newer Convex systems (such as the C3840; 960 MFLOPS peak) do not fair very well in his tests.

Himeno's experience with parallel machines certainly seems much more extensive than that at Toyota Motor or Toyota Central Research and Development Laboratory; it may be more extensive than that of anyone at a Japanese automobile company. However, we suspect that it is still somewhat lacking relative to, for example, the experience of a typical researcher at a U.S. national laboratory. During our visit Himeno mentioned that one of the central problems for researchers in Japan is access to advanced parallel machines for experimentation. He is one of several scientists we met during our tour who asked if access to parallel machines at Los Alamos through some kind of collaborative research would be possible.

We asked Himeno about how automobile companies manage to gain technological advantage over one another with computing. For example, we pointed out to him that all auto companies in the world have basically the same Cray computers and they all run crash worthiness simulations on them using, among others, the same "PAMCRASH" software. Himeno had two answers to this, the second of which was rather surprising. First, he said that companies differ significantly in the way they apply crash simulation software. For example, Toyota chooses to run their simulations using a finer computational mesh than does Nissan, but because of the extensive time that it takes to generate such a mesh, Toyota cannot carry out as many simulations as Nissan. More importantly, Himeno believes that advantage in high-performance computing is directly related to available CPU power. (For this reason, he looks jealously to the recent installation at Ford Motor Company of two CRAY C90s, the most powerful computing system by far at any automobile company, and indeed, one of the most powerful at any commercial entity in any field.)

Himeno also echoed a sentiment expressed during our previous visit to Toyota Central Research and Development Laboratory, Inc., on the general subject of computer simulation in research and design work. He said that there is still a relatively older generation of managers at Nissan who are somewhat more reluctant to rely on simulation than he would like, but as time goes on, this group of managers will be replaced by more computer-savvy ones who will look more to the simulation method.

3.7 CRC Research Institute, Inc.

CRC Research Institute was founded in 1958 as the Tokyo Electronic Computing Service Co., Ltd. whose primary business was the sale of G15 Bendix computers. When sales were not forthcoming, they began to sell cycles on their machines. This was successful and from that has evolved the current company. Reflecting this evolution of emphasis there have been several name changes over the years, until 1991, when the name was changed to its current one, CRC Research Institute. The company, with annual sales of about U.S. \$200 million, employs about 900 people, of whom 30% work in the Science and Engineering Group and 10% work in Research and Consulting, which is partly a marketing group. CRC is primarily a service company with little actual research support.

Our visit was to CRC's Makuhari Development Center located in a new "intelligent building" in the Makuhari New Tokyo City Center in Chiba. Our primary host was Kyukichi (Eugene) Ohmura, who is Managing Director of the Science and Engineering Group. In addition to his CRC-related activities, Ohmura has served as Asia/Pacific Regional Director of the Cray User Group.

Inquiring about CRC's customer base, we were told that CRC is an independent organization, meaning that they belong to no *keiretsu*. The advantage of this according to Ohmura, is that they can get business from everyone. Their business is 70% from private industry and 30% from government contracts. The government work comes partly from large construction projects such as bridges, dams, etc. and the industrial work largely from the nuclear power industry. There is a sizable component devoted to business applications, such as database maintenance and financial record-keeping; however, each individual job tends to be small.

They say they sell "complete solution systems", which means that when a customer comes to CRC with a problem, they provide problem setup, code development, and expertise. The customer then may purchase the solution code to run on their own machines, with continuing consulting from CRC, or let CRC maintain and run the code with differing input requirements.

Among the "complete solution systems" they provide are Engineering Systems, such as systems for nuclear power plants. The nuclear industry generates over 50% of their business, mostly in applied research. Dr. Hiroyuki Kadotani, who is Assistant to the President of CRC's Science and Engineering Group, does nuclear-related research work, primarily with power plant customers. He is familiar with the Los Alamos Monte Carlo neutron transport code, MCNP, as well as other publicly-available transport codes. He is very interested in using parallel computing in his research work, but says that he has no access to parallel machines. The Computational Science Workshop, run by LANL, interested him very much; he feels that attendance might give him the opportunity to learn about parallel processing.

CRC primarily uses codes that they have developed in-house, so that, once developed, they have the potential to be licensed and sold, in addition to being used for other customers' problems. Some of their

large established customers include the Japan Atomic Energy Research Institute (JAERI), NASDA, which is Japan's version of NASA, the National Highway Department, and Hitachi. With a total customer base in 1993 of 500, software portability becomes a problem, especially for licensing and selling software.

They also have collaborations in other countries, such as one with Century Dynamics Co. in Berkeley, CA, as well as something that they call "Overseas Business Partners". These "partners" include Cray Research, Inc. and SRI International of the U.S., UNIRAS A/S of Denmark, the Korea Advanced Institute of Science and Technology, and the China International Trust Investment Co., Ltd.. It is not clear what function these partners serve for CRC.

CRC has recently replaced its CRAY X-MP supercomputer with a CRAY Y-MP/EL, which is a deskside system equivalent in performance to a scientific workstation. Ohmura and Kadotani feel that, for their purposes, supercomputers are not needed. They feel that the current trend toward powerful workstations is much more compatible with their workload. They have several mainframes, however, made by Fujitsu, for their business applications and database work.

The economic hard times have had an effect on CRC and they see little chance for much growth in the near future. They feel that when growth comes, it will be in the area of multimedia applications and they hope to be positioned to take advantage of it through their software development program.

3.8 Taisei Corporation

Taisei Corporation is a large Japanese construction company founded in 1873. There are a total of about 14,000 employees, working in 12 Divisions and annual sales of about U.S. \$20 billion. Taisei is involved in all facets of large-scale construction projects from hotels and office complexes to bridges and dams. Their projects are located all over the world, including many in the United States.

We visited two parts of the company, the Technology Research Division, which is where the company's R&D is carried out, and the Information Systems Department, which houses and operates the company's important information processing equipment. There are about 400 employees in the Technology Research Center, which is in Yokohama, and the Information Systems Department, which is Tokyo, employs 140 people.

Our primary host at the Technology Research Center for this visit was Dr. Shunji Fujii, who is Chief Research Engineer in the Geotechnical Engineering Research Group. Fujii and his research colleagues use a Fujitsu VP2600 supercomputer, which is a 5-GFLOP system that is the top of the line in Fujitsu's vector supercomputer product line. Although the VP2600 is physically located in Tokyo, the Technology Research Division uses 80% of the available CPU time of the machine. A sophisticated network connects the research facility to the machine in Tokyo. Even with this sophisticated network (that includes a HIPPI interface), the ability to access the machines from home through modem dial-up is not available.

There were several presentations from scientists at Taisei. Mr. Shigehiro Sakamoto has been working for five years studying the fundamentals of wind flow around structures. He uses a grid size of 60,000 mesh points to simulate the action of wind around a structure shaped like a square cylinder. Previously, an older model, the "RANS" model, was used which gave time-averaged values of flows, information that is not useful for current wind studies. What are needed as input to the current model are time history values and the large Eddy simulation (LES) method gives these time-history values. LES is much more compute-intensive, however. It takes about 75 hours to simulate one case of an oscillating cylinder using the Marker-and-Cell (MAC) finite difference method to solve the Navier-Stokes equations. For a 400,000 grid-point job, the VP takes 270 hours. In order to complete this calculation, the job is run in two-hour segments. The square cylinder is the "worst case" simulation for wind action on very tall buildings. Mr. Sakamoto said that currently the VP2600 does not have enough power to support the fine mesh required to adequately represent the problem. Sakamoto, alone, uses about 65% of the CPU time available to Technology Research on the Fujitsu VP2600.

Dr. Yasushige Morikawa talked about numerical simulation of thermal and air flow distribution in large indoor spaces. His code models heat and air flow only. In modeling a large office building, he uses about a 34,000 grid-point mesh for office spaces, and for an atrium he uses different numbers of grid-points for modeling summer and winter flows. His "Thermal Canvas" program runs about 2 hours on the VP; it takes about 2-3 weeks to do the visualization of the resulting data on an SGI workstation. On the VP, the

rate for his code is about 500-600 MFLOPs. Dr. Morikawa is very interested in moving to MPPs, but he has no experience with parallel processing, a situation we found nearly everywhere we went. One reason for his interest is that he cannot model a whole building with the current technology. He needs either an MPP or something with 1,000 times the power of the current VP.

Taisei is also studying numerical methods in rock mechanics using a code they developed that does three-dimensional deformation analysis using the Finite Element Method. Results are used in the design and construction of tunnels. This analysis involves solving a large number of simultaneous coupled non-linear partial differential equations.

Other projects in the Research Division include hydraulic analysis, indoor thermal distribution, urban thermal distribution, and tsunami research.

One very impressive aspect of the Technology Research Center is the experimental facilities there. One is an apparatus called the 3-D Shaking Platform. Using this, scientists are able to simulate several seismic motions with a lateral force of up to 1G. By placing a scaled model of a building on this machine, they are able to determine the effects of various magnitudes of earthquakes on the structure. We also saw a facility that can simulate wave action on islands or shore facilities. It is located in a large building fitted with a wave machine to generate wave action. A large portion of the building can be flooded and instrumentation then measures the effects of this wave action on structures placed in or near the water. Dr. Fujii says that on hot summer days, they have been known to go for a swim in their "indoor sea." There is a wind tunnel that is used to study the effect of wind on various combinations of structures, such as those found in cities (tall buildings of various heights and large plazas, for instance). Maquettes are built and placed in the instrumented tunnel where wind velocities and effects can be measured at various places within the mocked-up structures and designs modified if necessary.

The availability of this extensive experimental equipment means that more weight is usually given to results from experiments than to results from computer simulations. But these experimental facilities are very expensive to build and maintain; therefore as simulation codes become more accurate and as older, less computer-sophisticated managers are replaced by younger managers more at home in a computerized world, reliance on simulation results is expected to gain acceptance.

As mentioned above, Taisei's Information Systems Department (ISD) which manages Taisei's computers and provides the infrastructure for computing throughout the entire Taisei operation, is located on several floors of a rather ordinary office building in the Shinjuku district of Tokyo. In addition to the VP2600, there is a Fujitsu M780 and an IBM 3090, both mainframes. The manager of the Information Systems Department is Mr. Masaaki Nakayama.

One of the most interesting aspects of Taisei is the network that is used to connect its supercomputer, its mainframes and the researchers at Taisei Research Division. It is one of the most sophisticated ones in Japan of which we are aware. The IBM 3090 serves both as a mainframe computer and as the front-end for the VP2600. These two machines share a disk and are connected by a high-speed channel. Access by

researchers is over Ethernet and a FDDI network from engineering workstations (EWS) and multifunctional terminals. An SGI Skywriter is connected to the VP by a HIPPI channel using Ultraset and is used for visualization.

The tasks of the Information Systems Department also include engineering analysis, graphics, structural analysis, on-site support to construction projects, development and support of business software, support of administration, system & network control, and information services, including support of databases.

This company makes heavy use of its supercomputer, although experimental results are preferable to simulation results which is understandable at this time, given the impressive collection of experimental equipment they have. They support construction sites through PCs on-site and at the ISD office. The network that Taisei has is the most advanced of any industrial site that we visited. Their use of computing in all phases of construction and design seems very advanced to us, but this may be more of a reflection of our naiveté about the construction industry than of their sophistication. They put forward the premise that the main reason for their use of supercomputers is that once one construction company in Japan buys a supercomputer, all construction companies have to have one.

3.9 Hitachi, Ltd. General Purpose Computer Division

In Japan there are several large companies whose business is referred to as "electric machinery manufacture." This term relates to a large number of products spanning the range from nuclear power plants to elevators to earth moving equipment to computers. The three largest industrial conglomerates in this sector are Hitachi, Toshiba, and Mitsubishi Electric. Of these, Hitachi is by far the largest, with annual sales for 1992-1993 more than 1.5 times that of the next largest competitor and profits for that period more than three times that of the next largest. Hitachi's annual sales have reached nearly U.S. \$80 billion, and its importance in Japan is suggested, in part, by the fact that it occupies positions in two of the seven large horizontal *keiretsu* groups that dominate Japanese business. In 1992, Business Week reported that Hitachi, alone, accounted for nearly 2% of the Japanese GNP.

In addition to its undisputed economic force, Hitachi, Ltd. occupies a unique position among its peers in the electric machinery industry because it is the only one that is also a major supplier of information and electronics equipment. In fact, Hitachi is one of the three manufacturers of supercomputers in Japan, and its latest machine, called the S3800, holds the current record for peak processing speed among vector supercomputers. The company has announced plans to build a large massively-parallel computing system which also promises to be one of the world's most powerful.

An important question is to what extent does Hitachi's prominence as a computer and supercomputer vendor affect its ability to improve its other products by using its own supercomputers for research and development. One source we have suggests that there are 10 Hitachi S820 supercomputers installed at various branches of the company. The S820 is the second-generation Hitachi supercomputer (the S3800 is the third), a vintage-1987 machine with a maximum of about 3 GFLOPS computing power.

In an attempt to explore how Hitachi uses its own supercomputers in its R&D we visited the General Purpose Computer Division (GPCD), which is located in the town of Hadano, in Kanagawa Prefecture. This location is also known as the Hitachi Kanagawa works, and employs over 3,000 people on a 52-acre site. All Hitachi computer systems larger than PCs are developed and manufactured here and at another site called the Hitachi Ebina works. Our visit was hosted by Dr. Shun Kawabe, Mr. Yoshiaki Kinoshita, Mr. Katsumi Takeda, and Mr. Hideo Wada, all of whom are involved with RISC processor development.

In order to describe to us the usage of Hitachi supercomputers within Hitachi, Ltd. proper, we were provided with a hand-out showing the "Supercomputer Network for Hitachi Laboratories." There are nine research laboratories within the company, and apparently supercomputers are also used at "production sites" but our discussion did not cover those. The nine research laboratories are: Central Research Laboratory, Advanced Research Laboratory, Production Engineering Laboratory, Image & Media System Laboratory, System Development Laboratory, Energy Research Laboratory, Hitachi Research Laboratory, Microelectronics Products Development Laboratory, and Mechanical Engineering Laboratory. Hitachi has recently (1991-1993) decided to consolidate the organization of these labs and

with this move to a centered system, supercomputing is also reorganized so that all of the labs share two main resources. At the Central Research Laboratory there is a HITAC S820/60 (1.5 GFLOPS peak) shared among five labs (Central, Advanced Research, Production Engineering, Image, and System Development) and at the Energy Research Lab there is an S3800/180 (8 GFLOPS peak) shared among three labs (Energy, Hitachi Research, and Mechanical Engineering). Additionally, some of the labs have machines of their own: Central Research has an S810 (.ca 1983 with 0.6 GFLOPS peak); Advanced Research has an S820/60 (1.5 GFLOPS peak); Energy also has an S820/60 and Mechanical Engineering has an S810.

Some of these labs have Hitachi scalar mainframe computers (M-680, M880, M682H, and M280H), some of which are very high-performance machines themselves, in addition to the vector supercomputers in use (S810, S820, and S3800). Nevertheless, what is clear is that there is not, as we might have expected, essentially "unlimited" supercomputing power installed within Hitachi, even though the company manufactures its own supercomputers. Our hosts reported that research budgets for the various Laboratories are managed independently, and that acquiring a supercomputer is treated much the same as acquiring any other advanced type of research tool, with similar budgetary constraints. Also, the majority of the supercomputers installed at Hitachi are from the older generation(s) of supercomputers, including two vintage-1983 machines that are still operating. More than one supercomputing observer has noted in the past that the older Hitachi machines generally performed less well on benchmarks than contemporary machines manufactured by, for example, Cray Research in the U.S., and Fujitsu and NEC in Japan. Finally, given that there are about 200 users of the machines, we would say that, in general, the Hitachi researchers are under-nourished with supercomputing sustenance.

One other interesting point is that all of the supercomputers installed within Hitachi run the proprietary Hitachi operating system, which is an IBM MVS clone. The S3800 is the only supercomputer Hitachi manufactures that is capable of running a UNIX-like system, but at Hitachi Energy Research Laboratory, even this machine runs the older system. Apparently, an adequate mechanism exists for transferring large quantities of data between the supercomputers running this system and UNIX-based graphics workstations (see below). However, our own qualitative experience suggests that the Hitachi operating system is quite cumbersome to use and lacks a great many features, even the simplest ones, of more modern operating systems. Clearly, Hitachi engineers and scientists are able to produce results on their machines; however, we would estimate that in general, productivity is lower than it might be.

All of the machines described above are interconnected via networks that run at relatively low speeds. Within each of the two consolidated laboratories, there are interconnections with speeds listed as 3 megabits/second, 384 kilobits (kb)/second, and 192kb/second. The connection between the consolidated labs themselves is listed currently at 384 kb/second and an upgrade to twice that speed is planned.

While at Hitachi, we were shown a video with several examples of supercomputer simulations that have been carried out at various company laboratories. Included were:

- Air temperature and air flow simulation (about 5 CPU minutes) in a car of the *Shinkansen* (bullet train)
- 3-D Finite Element Method electromagnetic simulation of the superconducting track for a magnetic levitation train showing the Eddy current distribution, from a simulation that required about 10 minutes of CPU time.
- Unsteady flow in a turbine stage rotating at 132 and 401 meters/second using the k-e (non-turbulent) Finite Volume method.
- Propagation of a Tsunami (tidal wave) caused by an off-shore event such as an earthquake. The simulation used a shallow-water wave model on a 400 X 450 grid and required about 3 CPU minutes to run.
- Turbulent fluid flow around a cylinder with Reynolds number $\sim 10^5$
- Molecular Dynamics simulation to explore both ductile and brittle fracture formation in an iron slab. The simulation used 1,000 particles for 3,000 timesteps and was run on a HITAC S820.
- Combustion simulation that required about 3 hours of computational time using a very small grid (35 X 55).
- 2 -D turbulent flow in a variety of simple geometries such as T-joint, Y-joint, and expanding tube with a perpendicular bend. Simulations ran for about 40 minutes.

The visualization effects in all of these simulations were excellent. They had been carried out largely on Silicon Graphics "Iris" workstations. In general, the simulation results seemed very impressive, both in terms of the computation itself and the visualization. However, with one or two possible exceptions (the trains and the turbine), it was difficult to see how the simulations related to actual products. We were unable to get answers to this question.

Separate from the Laboratories is the General Purpose Computer Division itself, which apparently has six HITAC S820 supercomputers of its own. Undoubtedly, some of these are used for system software development, as is the case with most other computer vendors. However, some are used for logic and device simulations for new Hitachi computing systems. In particular, some of the supercomputers have been modified specifically to include special-purpose hardware and new instructions that allow logic simulations to be carried out at much higher processing rates than is possible with conventional systems. Hitachi has developed (and published papers on) VELVET, which is a vectorized, event-driven, gate-level and register-level logic simulator that supposedly runs two orders of magnitude (100 times) faster than "conventional" (software-based) gate-level simulators. VELVET was initially implemented on an S810 supercomputer equipped with this new hardware and then ported to the newer S820 model that was similarly equipped, but has not yet appeared in the S3800. Hitachi claimed that unless there is outside interest in having this special-purpose hardware on the S3800, there are no plans to make it available. Each of the six new hardware instructions implemented in VELVET can execute the logical operation of a

gate being simulated in a single S810 or S820 clock cycle. A paper published by Hitachi personnel at the ACM/IEEE Design Automation Conference in 1988 states that VELVET running on an S810 was able to carry out, in a reasonable period of time (about 100 minutes), an entire-machine simulation of the Hitachi S820 supercomputer. Our Hitachi hosts claimed further, that their newest supercomputer, the S3800, which, by the way, has the fastest executing silicon gate arrays on any computer available in the world, more than twice as fast as those from Cray Research, could not have been developed without using VELVET. The system is also used in developing RISC-based microprocessors. Hitachi is the major supplier of Hewlett Packard's latest RISC microprocessor, the PA-RISC, which is used in scientific workstations. Also, Hitachi will use VELVET in developing its own microprocessors, such as the ones that will be included in their future massively-parallel system, the CP-PACS.

There are two other systems similar to VELVET in existence, computing systems that were designed specifically for the purpose of accelerating logic simulations. One is the Yorktown Simulation Engine, a 256-processor machine used by IBM and the other is HAL2, a 64-processor system developed by NEC. On the other hand, we learned recently that Cray Research uses a standard vector processor for its logic simulations. We wanted to compare the relative speeds of Cray's simulator and Hitachi's VELVET system, so we asked Tom Court and Steve Oberlin who are designers at CRI for help. They estimate that Cray's simulator running on its standard Y-MP supercomputer is still about 8 times faster than the special-purpose Hitachi S810 system. And, while Hitachi says it will not port VELVET to its newest supercomputer, the S3600 or S3800, Cray will have no problem porting its simulator from the CRAY Y-MP to the CRAY C90 and other vector-parallel machines that will succeed the C90.

3.10 Angstrom Technology Partnership's Joint Research Center for Atom Technology and Agency of Industrial Science's Research Information Processing Center

MITI, the Ministry of International Trade and Industry, through its Agency of Industrial Science and Technology (AIST), has launched a government effort to pursue research and development for "precisely observing and manipulating individual atoms and molecules, on a surface or in free space, and to its supporting technology" [1] through a project titled Ultimate Manipulation of Atoms and Molecules (Atom Technology). The National Institute for Advanced Interdisciplinary Research (NAIR), formed in January 1993, provides the infrastructure for government support of this ambitious initiative. Objectives of the ten-year project include observation and manipulation of individual molecules and atoms on solid surfaces and in space (nanotechnology), creation of technologies to develop atomic level structures, and theoretical simulation of atomic and molecular processes. The project hopes to achieve results by the year 2001 and to disseminate its findings through international symposia, extensive publications, and through the employment of postdoctoral fellows. The project is budgeted for 25 billion yen (\$210M) over ten years. This research is not being implemented in a decentralized fashion, but rather, all researchers are brought together in Tsukuba to do their work. For the first six years, R&D will be pursued on basic and enabling technologies; the second phase will concentrate its efforts in developing new materials, electronics, chemicals, and biotechnologies.

The Angstrom Technology Partnership (ATP) is a consortium of approximately 30 companies established in February 1993 as a result of this research endeavor. The companies collaborating on the project are international in nature and span such technological areas as computing, chemical industries, heavy industry, biology, and electronics. The work is carried out at the Joint Research Center for Atom Technology (JRCAT), which maintains a well-endowed, state-of-the-art computing facility.

Dr. Satoshi Sekiguchi, who is in the Computer Architecture Section of another AIST institute, the Electrotechnical Laboratory (ETL), arranged our visit to these government laboratories. At JRCAT we spoke extensively to Mr. Katsumi Nakayama and Mr. Koichi Sato of the systems section at JRCAT. Proving once again that computer system managers the world over are cut from the same mold, Sato and Nakayama were enthusiastic about their jobs and equipment, eager to talk to us about it, and appeared very competent. JRCAT has a Thinking Machines Corporation CM-5E and a Fujitsu VPP500 computer, both of which are new parallel processing computers. The CM-5E has 128 processors, with a peak speed of 20 GFLOPS, 16 GB of memory and 44 GB of SDA storage (disk). On a molecular dynamics code they have recently run, scientists have achieved a performance of 6.5 GFLOPS using 64 processors; they are hoping to obtain 10 GFLOPS using the entire capacity of the machine, 128 processors.

ATP's VPP500 is a 32-processor machine, with 256 MB of memory per processor, 44 GB of disk capacity, and an impressive 51.2-GFLOP peak speed for this configuration. Dr. Takehide Miyazaki, a research scientist at JRCAT, spoke with us regarding his molecular dynamics application. Dr. Miyazaki is studying the process of crystal growth on silicon (Si) by molecular beam epitaxy using molecular

dynamics simulations. This work has direct benefits to the industrial sponsors of JRCAT such as NEC, Fujitsu, and other companies involved in semiconductors. He is one of the first users of the VPP and has achieved a 4-fold speedup using 7 processors. His code, consisting of about a thousand lines, has been adapted for parallel execution on the VPP500 with the addition of compiler directives. However, Miyazaki said that he normally runs the code only on a single processor and that the majority of people are currently using the machine in this manner.

A short distance away from JRCAT is the computer facility associated with a facility known as Research Information Processing Center (RIPS). RIPS is the computer technology center which the nine Agency of Industrial Science and Technology laboratories at Tsukuba use for solution of their research problems. This center, which networks all the laboratories via fiber optics, was initiated in 1981 with a Fujitsu FACOM M-200 mainframe computer. In 1984 the facility was upgraded to house two mainframe computers, a FACOM M-380 and an IBM 3081K. At the same time, a network was put into place which linked all seven of the AIST laboratories in Japan. A supercomputer was first installed in 1987, at which time the local network was enhanced by the addition of a high-speed channel. The facility, like many that we saw in Japan, has earthquake safeguards installed, as well as an energy-saving control system and protection against fires.

RIPS maintains the following high-performance computing systems:

- Cray Research Inc. C90 - 4.2 ns clock; 1024 MW total memory., 1024 MW DRAM SSD (extended storage), cooled by fluorinert. The RIPS CRAY C90 has 16 processors, which is the largest C90 configuration available.
- IBM Power cluster - 16-processor cluster of RISC System/6000-590 workstations for scientific applications.
- DEC 10000 - an alpha chip-based, two-CPU machine, with 250 GB capacity, which is used as a file server.
- 2 Cray Research Inc. CRAY 6400s (Model EL) - each machine has 32 processors, is SPARC-based, uses NQS to facilitate dispatching jobs to the C90.

For networking purposes, there is a HIPPI interface between the C90, DEC 10000, and IBM cluster machines.

This was the first facility we had seen that had a significant number of non-Japanese computers. Overall, the amount of computing power in both computer rooms was extremely notable and the persons we spoke to appeared quite literate in computational jargon, at the very least.

3.11 Grumman Data Systems Corporation

A subsidiary of Grumman Corporation, Grumman Data Systems Corporation was founded in 1969 as a computer software and service bureau. The company also does custom-designed integrated information processing systems and large-scale computer systems integration. There are currently about 2100 employees. The Data Center is managed by a president and 8 vice-presidents. The parent company has recently been bought by Northrup Corporation, so it is unclear at this point exactly what the future of the Data Center will be. Grumman, the parent company, was comprised of several companies: Grumman Aerospace, Grumman Data Systems and Support, to name two.

In late May of this year, Grumman Data Systems shut down its CRAY Y-MP/2E. This machine had been used primarily for customers within the Grumman umbrella, mostly Grumman Aerospace. The Aerospace group came to the conclusion that they could handle most of their applications with powerful workstations and buy supercomputer time more cheaply elsewhere than through Grumman.

One of Grumman's efforts is as a system integrator for other companies. In this capacity, they have installed and manage several Cray sites, including the Naval Oceanographic Office in Mississippi. It appears that these Cray sites are also being phased out, at least where Grumman is concerned. It would appear that at least for the near future, computing services at Grumman will likely shift emphasis toward the use of workstations.

3.12 National Center for Supercomputer Applications

The National Center for Supercomputer Applications (NCSA), located on the campus of the University of Illinois at Urbana-Champaign, was formed to provide supercomputing access to academic and industrial researchers by adapting the best technologies and concepts from the national laboratories and replicating these in a more accessible environment. Having in large part, declared victory on that front, the focus has shifted to providing more advanced, less expensive resources to a smaller community of users trying to solve grand challenge-type problems. Additionally, efforts are being made in the area of information technologies, cyberspace tools, and K-12 education.

Emphasis is on quality computing cycles and consulting in their use. Much effort is also going into providing public domain software for the masses. Customers of NCSA include academic and industrial researchers in need of computer resources and users of public domain software, especially that developed at NCSA, such as NCSA MOSAIC. Current industrial partners with NCSA are the following: Dow Chemical Co., Caterpillar Inc., FMC Corporation, Eli Lilly and Co., Phillips Petroleum Co., JP Morgan, Schlumberger, AT&T, Eastman Kodak, and Motorola.

About 200 persons are employed at NCSA, plus another 50-100 students, who work varying amounts of time. The staff is largely computer science-oriented; however, there are a handful of scientific researchers, primarily from physics, biology, and chemistry.

NCSA is owned by the University of Illinois but is funded largely by the National Science Foundation (NSF). Management structure consists of a Director, Deputy Director, and eight Associate Directors. Current hardware consists of:

- 512 -node CM-5 with 16 GBytes of memory, 1 HIPPI adapter, multiple FDDI and ethernet, 150 GBytes of SDA;
- 16-processor SGI Power Challenge with 2 GBytes of memory, 1 HIPPI adapter, 1 FDDI, 80 GBytes of disk array;
- hypernode Convex Exemplar with 512 MBytes of memory, FDDI, and 4 GBytes of disk storage;
- 2-processor Convex C220 with Unitree;
- 8-processor Convex C3880 (this machine is about to be removed);
- 4-processor CRAY Y-MP (this machine is about to be removed);
- 4-processor CRAY-2 (in process of removal);
- several hundred workstations (Sun, Apple, SGI, HP, DEC, IBM).

NCSA considers its chief competition to be the other NSF supercomputer centers, the state supercomputing centers, and the national laboratories. They feel that any edge they have over their competition has come, and will continue to do so, by simply concentrating on being better at providing their services. Software for a wide variety of applications is both commercial and produced in-house.

Much optimization is done on software. The current customer base is about 3000 users in the computing arena and growing smaller, while the public domain software effort in the center is growing by leaps and bounds. The majority of the applications run at NCSA are fairly small, and it is felt that the customers are beginning to feel that they can be better served by purchasing powerful workstations and doing the work themselves. Thus the emphasis has shifted to a smaller set of users executing very large problems.

Much work is going on at NCSA in the area of visualization and virtual reality. This is where many feel the strength lies at this institution.

3.13 Bechtel Corporation

Bechtel is a San Francisco-based global engineering and construction company that was founded in 1898 in the then-territory of Oklahoma as a builder of railroads. Success in railroad work led to general construction projects that include dams, bridges, power plants, space facilities, pipelines, and other related projects. Two of their latest projects involve building an entire city, Jubail, in the Arabian desert, and building the Channel Tunnel between France and Britain. Bechtel employs a total of 20,000 people worldwide. In 1993, they reported revenue of approximately \$7.3 billion, representing work from about 600 clients. Due to time constraints, information on Bechtel's computer usage was obtained through telephone conversations, written material, and electronic mail exchanges.

Like many of its industrial counterparts in the U.S., Bechtel has used mainframe computers to do its business processing and any research computing for more than 25 years. This usage has not led them to purchase supercomputers for the research and simulation studies. With the advent of affordable scientific workstations with acceptable performance, most of their research work has moved from the mainframes to these platforms. Some simulation work has even been moved to PC-class machines now that their performance has reached reasonable levels for mathematical computations. These simulations are primarily related to industrial process and airport traffic studies. All of their research activities, in fact, are directly related to engineering and construction; they do no work on fundamental physics problems. This contrasts sharply with the Taisei Construction Company that we visited in Japan, whose research activities are also centered around engineering questions, but whose researchers also investigate fundamental physics questions. Instead of doing fundamental work in-house, Bechtel relies on frequent interaction with researchers in universities and National Laboratories for studies and computer programs related to advanced methods of analysis and modeling.

According to our sources, the justification for buying computers, in this case scientific workstations and PCs, and the measure of the benefits from their use, is based primarily on cost savings. These sources can cite no specific examples of reductions in research time or cost due to the use of computers, nor can they find examples of gaining knowledge through computer usage that could not have been gained by other means. The simulations that they do conduct have no effect on corporate decisions, which is much the same as we heard from many Japanese industrial researchers, even those who have large supercomputer operations and use them extensively.

One conclusion that one can draw from this, is that while U.S. construction companies do not use high-performance computers in their work, unlike their counterparts in Japan, neither country's construction companies base corporate decisions on the results obtained through computer simulations, no matter what the platform. Bechtel has an alliance with several Japanese construction companies, and we were not able to find out if these alliances include sharing of research and simulation results.

3.14 Los Alamos National Laboratory Engineering Sciences and Applications Division

The Engineering Analysis Office of LANL's ESA Division has several scientists (two full-time Staff Members, one Postdoc, and several students) who use computers to study a variety of civil engineering problems. We spoke to one of them, Dr. Charles Farrar, to further understand the role of supercomputing (if any) within the construction industry in the United States.

Dr. Farrar and his colleagues carry out engineering analyses related to seismic stability for new buildings at LANL. Dr. Farrar is also well known within the nuclear community and has carried out similar assessments of power plants for the U.S. Nuclear Regulatory Commission. Their team has funding from many external sources for other analyses, such as highway bridge evaluation. They also conduct some basic research related to geology of faults, trying to understand the mechanism of energy release during earthquakes.

The assessment of both bridges and power facilities centers on analysis of deformation damage caused by vibration. This is usually estimated through experimentation, and Farrar's group carries out such studies even though they are often very expensive. However, computer simulations using the finite-element method are also an important means of evaluation. The simulations can predict both the "shape" of the deformation and also what conditions will lead to stresses that exceed the limits of the materials used. (Interestingly, Farrar mentioned that for many structures these limits are provided by the American Society of Engineers in the form of stress tables. However, the data in the tables are generally obtained from laboratory experiments, and up to now, no one has bothered or managed to verify them using a supercomputer. Farrar believes that there is a need to do this.)

Farrar mentioned that many civil engineering simulations are not complex enough to need supercomputers but others most definitely are. For the latter, Farrar uses LANL's CRAY Y-MP supercomputers. Some problems require that a large, sparse matrix be inverted up to 10,000 times, which takes many hours on a workstation, but can be done in a few minutes on the Y-MP.

Farrar also gave us several clear examples of how advanced computers can improve civil engineering simulations. For a given structure, sometimes the difference between a supercomputer simulation and a workstation simulation means that the response of a structure can be estimated for a longer period of time. For example, Farrar showed us simulations of how reactor containment vessels respond when subjected to an earthquake. As input data they use actual traces from seismic recorders during real earthquakes. The data trace can be up to 60 seconds in duration but Farrar's simulations only analyzed about 10 seconds of the trace due to the long simulation time.

Another example from his reactor studies concerned how geometric or material irregularities in a reactor vessel can require significantly longer simulations and yield significantly different results. Typically, a reactor vessel is highly symmetric, so that only a portion of it, say one-half, needs to be simulated.

However, a small defect at a single point can destroy the symmetry requiring that the whole vessel be evaluated to understand the effect of a defect in one part of the container. This adds significantly to the simulation time and can yield significantly different results from the symmetric simulations. However, Farrar also pointed out that relatively poor understanding of the physics involved may have an even larger effect. For example, how to include the effect of dampening of seismic waves is poorly understood. Dampening in Farrar's supercomputer simulations can easily mean the difference between a reactor that will collapse and one that will survive.

The finite-element computer code used in Farrar's team is ABAQUS, which is a widely-used third-party program from Hibbett, Carlson and Sorenson in Providence, Rhode Island. ABAQUS, and other programs like it, currently run on PCs, workstations, mainframes, and vector supercomputers, although supercomputer users are probably the smallest contingent. Notably missing from this list of computers are massively-parallel machines. Commercial products such as ABAQUS have not yet been ported to these machines.

Farrar mentioned that most construction companies in the U.S. do not do extensive finite-element modeling because they do not have access to enough compute power to be able to run these simulations economically. The engineering firms that do run simulations tend to use very simple models that can be run on workstations. Other work is done at various universities, such as the study of earthquake effects on highways and long elevated bridges being carried out at several California universities and funded by CALTRANS. According to Farrar, U.S. companies do not conduct as much testing in general as Japanese companies do and the Japanese construction industry has far better experimental facilities than are available in the United States.

A significant problem in civil engineering that would lend itself well to supercomputing simulations, Farrar feels, is something called "bridge scour." This occurs during flood conditions when the rushing water "scours" out the base around bridge piers, exposing the underlying pilings and endangering the bridge structure. Not much is known about how this occurs. He feels that computer modeling of this phenomenon is badly needed, and it is definitely a supercomputer problem because of the need to couple the simulation of the water flow with the stability of the bridge. So far, laboratory experiments at Colorado State University's Hydraulics Laboratory constitute the only work done in this area. Farrar feels that a lack of funding has prevented simulation efforts. It seems that most of this type of work is funded by the Federal Highway Administration and there appears to be a lack of eagerness there to change old ways of looking at problems and to embrace "new technology." However, since using the "old ways" leads designers to use overly conservative designs (according to Farrar), use of more sophisticated analyses and techniques through computation could save money.

Thus, perhaps the usage of supercomputers in civil engineering is not so different from that in other areas, in that not only are they useful research engines, but they can provide cost savings as well. Given that there are an estimated 500,000 water-based bridges in the United States, the potential savings could be quite significant.

3.15 Los Alamos National Laboratory Advanced Computing Laboratory

Several factors are involved in a comparison of U. S. vs. Japanese industrial supercomputer usage. Both organizationally and structurally, the two countries have approached the issue of high performance computing in a different fashion. We saw many supercomputers bought by Japanese companies; there appears to be a more modulated and cautious approach involving less risk in the United States. Through its numerous national laboratories and collegiate supercomputing centers, the country has provided a focus point for industries to "get their feet wet," as it were, prior to making a more significant investment of purchasing supercomputing hardware. National Laboratories that have to this point turned their vast research resources towards defense activities, are now in the process of redefining their missions. For many of them, including Los Alamos National Laboratory, this has included providing strong ties to related civilian industries, such as automotive, gas and oil, and financial. There is a strong push for technology transfer of skills and products between the laboratory and these companies. There is an especially active approach, in part a survival effort, to find and nurture these new "customers."

Concrete structures are in place to enable this to happen. The government has instituted formal Collaborative Research and Development Agreements (CRADAs) for projects of mutual benefit to industry and government. The process has been established to permit these collaborations to proceed in a timely fashion and is designed to cut through the bureaucracy that has traditionally hampered these activities. (Currently, at Los Alamos, about 150 CRADAs are in place). To our knowledge, such collaborations, at least through this formal type of structure, do not exist in Japan. One of the industries that is a significant participant in CRADAs is the petroleum industry. As there exists no counterpart industry in Japan, it is difficult to make a comparison of computer access for this sector in the two countries.

The Advanced Computing Laboratory (ACL), located at Los Alamos National Laboratory, was established in 1988 as an experimental facility to investigate high performance computing in an environment that would allow innovative, promising computing resources to be made available, without impacting the operations of the centralized production computing facility in Los Alamos. Researchers, both academic and industrial, as well as Laboratory scientists, are encouraged to pursue research on leading-edge computers. Because of the potential high performance of these parallel computers and their large purchase cost, it is often desirable, but unfeasible for researchers to test applications on them in their own environments. The ACL enables users to make a modest investment with limited risk before selecting (if at all) a computer for their own site. Payment for resources has been handled in a variety of ways, from contributing to the purchase of facility equipment to participating in an approved research program for which funding has been allocated to the Advanced Computing Laboratory. As time has gone on, many industrial concerns have become clients at the ACL, using Laboratory expertise to help solve their particular application and/or availing themselves of the massively parallel supercomputers in the facility. Currently, the ACL houses a 1024-processor CM-5 from Thinking Machines Corporation

(TMC) and a Cray Research, Inc. 128-processor T3D, as well as numerous high-end scientific workstations.

Because of the interest of industry in the program, a new Center called the Computational Testbed for Industry (CTI), was established and resides within the ACL. Established in 1993, the CTI offers several categories of membership, starting at \$10,000, in which a business may purchase both computing resources on the TMC CM-5, various Cray computers, cluster computing systems, and also Laboratory expertise on problem-solving through team efforts and on-site collaborations. The membership also provides for consulting and attendance at workshops on topics related to high performance computing. These agreements, called User Facility Agreements, number about 25 at the time of this report.

The presence and encouragement of a facility such as the ACL provides sharp contrast to the apparent lack of access to parallel and novel computing architectures that we encountered in Japan. An exception to this is the collection of research laboratories located in Tsukuba. As mentioned elsewhere in this report, parallel computers are installed in several sites there and appear to be functioning in the above fashion. Thus, the comparison is oblique because different environments and emphases seem to affect the whole picture. Other factors, as discussed in this report, must be interleaved with the above observations before any definitive conclusions are made about the relative strengths of the two countries in high performance computing.

3.16 Some Conclusions from the Case Studies

Though we were unable to visit companies in all the fields for which we originally had hoped, there are still many conclusions that can be made from the visits. Some can be drawn regarding the Japanese companies themselves and some can be concluded by comparing them to their U. S. counterparts, namely in the areas of automotive, service industries, and government research laboratories.

First, a wide variety of industrial users in both countries have employed some sort of high performance computing to achieve their product goals. The automobile industry in both countries relies heavily on supercomputers and the role of the supercomputer in this industry is fairly well understood. In other industries, such as the construction business, more initiative in supercomputing use has been shown in Japan. One reason may be that most U.S. construction companies are much smaller than the ones in Japan that own supercomputers (although Bechtel's size is certainly in this range and they don't use supercomputers). The Japanese use of supercomputing in the construction industry appeared to be a case of "follow the leader." As one company purchased a machine, others seemed to be eager to do likewise, so as not to be left behind. (We were told more than once that this is a common phenomenon in Japan.)

The use of supercomputing in chemical and pharmaceutical firms differs in both countries. Whereas companies such as DuPont are using their computing resources to do actual process control, most Japanese chemical companies are pursuing purely research issues on their supercomputers.

In both countries business at computer service industries appears to be declining, as many more clients are turning to less expensive, high-performance scientific workstations to solve their problems. Consulting on solution mechanisms and suitable algorithms may continue for them, and this may be the biggest draw of these companies in the future.

One area in which we saw an apparent difference between the U. S. and Japan is in the use of massively parallel computers. Whether this is because of the sluggish Japanese economy that is hindering research in this area or whether it is because of other factors was not always clear. Currently there is heavy reliance on vector supercomputers throughout much of Japanese industry. The automotive industry stressed the importance of simulation techniques to automobile development and is a large user of Cray-type vector supercomputers. At both Toyota Central Research and Development Laboratory and Nissan Motor Corporation, we did hear that they would like to have more computational power for computational fluid dynamics applications and were, therefore, interested in using parallel computers. It is uncertain whether this will come about in the near future.

Of the two companies we visited in Japan that make consumer products, Matsushita was the only one that we saw using supercomputing to develop a consumer product (rice cooker and microwave oven). Hitachi probably uses some of its older supercomputers in the development of its consumer products, but we were not able to obtain any information as to what these were. The Matsushita example was treated

with some chuckles when we mentioned it at other companies during the rest of the trip. Since rice cookers are big business in Japan, it seemed reasonable to us that heat studies and conduction flows would be modeled, but other companies inferred that perhaps Matsushita just "didn't know what to do with their supercomputer."

Although we saw several instances of supercomputers being used in product development, we also saw many examples of usage more related to basic research, such as at Matsushita, CTI, Toyota Central Research and Development Labs, and Hitachi.

We are aware of only one other attempt at a broad survey of computational science in Japanese companies. Focardi [1] visited Konoike Construction, NEC, and Ohbayashi, as well as a Japanese National Laboratory, five U.S. companies, a U.S. university, and a U.S. National Laboratory. Although a direct comparison of U.S. and Japanese supercomputer usage was not a specific objective of Focardi's work, he nevertheless drew several conclusions about differences he observed. Focardi carried out his research under contract for the Swiss National Government, and in particular, a Swiss national computing facility attempting to persuade European companies to form cost-sharing partnerships with it. Generally speaking, we believe Focardi's report portrays both computational science and Japanese R&D in a somewhat optimistic and flattering way. Shown below are two of Focardi's findings, and some brief notes comparing them with those of our own.

- "Japanese do not go through a cost justification process of investment to the same extent as their American counterparts." We did not find this to be especially true at the companies we visited. Perhaps Focardi's research reflected more the conditions during the Japanese "bubble" economy.
- "Top Japanese management is persuaded that investing in R&D [and therefore supercomputers] as a way of augmenting the company's assets. . ." In Chapter 2 we noted several examples of companies for which this was true.

4. Summary and Conclusions

The primary goal of this report was to compare supercomputing usage in the U.S. and Japan. In contrast with last year's study [1], which focused on supercomputer manufacturing capabilities, the primary focus of this work is to compare where, and how well supercomputers are being used in the two countries.

There were two important motivations for carrying out such a study, as mentioned in the introduction. The first was our prior impression that more Japanese companies used supercomputers than American companies. However, the survey data in Chapter 2 of this report show that this is not the case, if the most powerful supercomputers that exist today are the basis for comparison. Our previous impression had been based on data collections for Japan that included many machines now considered to be low-end models. Today, roughly the same number of businesses in the U.S. and Japan are using supercomputers.

The difference we see between the current tallying and the previous studies means that many Japanese companies are not purchasing upgrades to their older systems. They are apparently "making do" with less performance than they could get if they bought a machine today. Given the importance of supercomputers in shortening the development cycle for commercial products, we believe they are therefore at a disadvantage. We postulate that any simulation that can be done on a supercomputer can always make use of more computational power. Either longer times can be simulated or a finer scale can be used or the effect of a greater number of parameters can be studied. We doubt that they have given up on the idea of supercomputing altogether. More likely possibilities are (1) economic conditions prohibit the expenditure; (2) they are reluctant to try parallel machines (see below); (3) they are waiting for Japanese manufacturers to provide parallel systems instead of having to buy one from an American vendor; and (4) they are using high-performance workstations instead of supercomputers.

A second important motivating factor in this study was the fact that Japanese supercomputers are equipped with user environments that match that of Japanese mainframe computers. This compatibility might have been why more Japanese companies use supercomputers than their American counterparts, because it would have been easier for Japanese companies to upgrade to supercomputer-class machines. Although this was an advantage in the earliest years of supercomputing in Japan, our research for this report suggests that both Japanese supercomputer vendors and Japanese supercomputer users actually suffered a detriment by having mainframe-compatible supercomputers. This is because in the late 1980s, the supercomputing world in the U.S. and in Europe evolved to "open" supercomputing environments that offered greater software and interconnection capabilities. However, with their own proprietary operating systems, not only did Japanese vendors lose sales, but companies using Japanese supercomputers also became isolated, lacking in applications software, and probably suffering an overall loss of productivity in their supercomputing research. Although this situation is now improving, as all three main Japanese vendors move to more open, UNIX-like systems, our field research (Chapter 3)

suggests that many users of Japanese supercomputers are still not operating in what we regard as state-of-the-art, "user-friendly," productivity-enhancing environments.

Recently, one of the authors attended an external review of the Los Alamos Computing, Information, and Communications Division, at which it was asked, "who are the main competitors of Los Alamos for the title of 'World's Best Computing Facility'." Among the names mentioned were the National Center for Supercomputing Applications in Illinois, the DOD High-Performance Computing Center in Minnesota, and Japan's National Aerospace Laboratory. The last of these was included because it currently houses what may be the world's most powerful supercomputer, the NWT machine made by Fujitsu. However, while the best hardware is certainly an important component of the resources needed to apply supercomputing, other factors may be equally as important.

In Table 4.1 we summarize what we believe to be the relative status of the U.S., Japan, and Europe in several of these other factors, what have been called "core subfields" of computational science [2]. For this subjective comparison we use a scale that covers the range - (worst), 0, +, ++ (best). An arrow next to a rating indicates the expected change in that rating over the near term. The "data" for this table come from our own various impressions: from the supercomputing literature; from using supercomputers, and from our field research.

The United States currently leads Japan in all areas except hardware reliability, which has been a problem for some recent high-end supercomputers. In several key areas, especially high-speed networking and multiple-processor design, Japanese technology is behind but rapidly improving. In both of these areas Japanese companies possess many core strengths (especially in electronics, optics, and semiconductor device fabrication) that may allow it to catch up soon or even catapult into the lead. In other areas in which the U.S. leads, such as user interfaces, the advantage is unlikely to last, because such technologies are easily duplicated or are readily and freely available anyway. Of course, Japanese users will always have somewhat more difficulty with interfaces because of the complexity of the Japanese written language.

Table 4.1: Comparison of the United States, Japan, and Europe in Several Important Computational Science Disciplines.

Core sub field	U.S.	Japan	Europe
multiple processors:			
design & manufacture	+	0 ↑	-
use	+	- ↑	+
Data Communications and Networking	++	0 ↑	0 ↑
Software Engineering	+	+	0
Information Storage and Management	+	0	0
Hardware Reliability	-	+	N/A
User Interfaces	+	0	+

In comparing U.S. and Japanese supercomputing usage sector by sector, we found, not surprisingly, that supercomputing activity reflects the overall commercial strengths in each country. For example, much of the United States' industrial supercomputing power is in the petroleum industry. In contrast, in Japan there are many supercomputers being used in manufacturing and electronics companies and in the U.S. there are very few.

In Japan there are also a few relatively novel uses of supercomputing within industries, notably the construction industry. Although construction industry supercomputing doesn't seem to have much effect on construction industry corporate profits, this seems to be an example of how Japanese companies are willing to make large investments in capital and personnel in order to pursue basic research and to "modernize" by encouraging use of new, high-technology solutions. There are other examples of Japanese companies that also took this same optimistic, exploratory attitude.

One sector in which supercomputing usage in Japan is increasing substantially is the aerospace industry. In both private and government facilities, but especially the latter, supercomputing is viewed as a critical portion of an industry the Japanese government has specifically targeted for extensive growth. Supercomputers are being used in Japan in the design of both domestic and foreign commercial aircraft, and in the design of what Japan hopes will be its version of a space shuttle. What very well may be the world's most powerful supercomputing facility is at the Japanese National Aerospace Laboratory. In contrast, note that several commercial aerospace companies in the U.S. are currently downsizing their supercomputer operations, due to lack of need for supercomputer resources.

In the U.S. the government's role as primary funding source for supercomputing is still quite considerable. If one counts both government labs and government-funded supercomputer installations at universities, the U.S. Government is still the largest funding source for supercomputing. In Japan the government's role is growing substantially, and over the last two years or so, government supercomputer purchases have far over-shadowed those of Japanese industries, both in quantity and value. The Japanese Government's "economic stimulus" program, the source of most of the supercomputer purchases, is a much larger form of support for the supercomputer industry than Japan has undertaken until now. The program has primarily benefited Fujitsu, which is currently the strongest of the big three Japanese supercomputer vendors, both in terms of technology and supercomputer profitability.

In addition to the function as an economic stimulus, Japanese Government supercomputer purchases are intended to improve national research facilities that had been under-funded in computational capabilities and to improve basic research in Japan. Although Japan has not explicitly identified "grand challenge" computing applications, an estimate of what they might be can be made by noting where the supercomputers are being placed. A comparison of U.S. and Japanese "grand challenges" is shown below [3]. For the Japanese entries, the installations at which most of the work is being carried out are given in parenthesis.

Grand Challenge Computing Applications in the U.S.

- Aircraft
- Combustion Modeling
- Particle Physics
- Environmental Modeling
- Molecular Biology and Biomedical Imaging
- Product Design and Process Optimization
- Computing Education
- Plasma Physics
- Oil Reservoir Modeling
- Space Science
- Parallel I/O

Possible Grand Challenge Computing Applications in Japan

- Aircraft (NAL)
- Flexible Computing (RWC)
- Solid State Physics (RIKEN, NRIM)
- Nuclear Power (JAERI and PNC)
- Single-Atom and -Molecule Properties (JRCAT)
- Telecommunications

Abbreviations:

NAL: National Aerospace Laboratory

RWC: Real World Computing Program

RIKEN: The Institute of Physical and Chemical Research

NRIM: National Institute for Metals Research

JAERI: Japan Atomic Energy Research Institute

PNC: Power Reactor and Nuclear Fuel Development Corp.

JRCAT: Joint Research Center for Atomic Research

The United States is now well ahead of Japan in beginning to apply parallel processing to scientific and engineering problems. Hardly any Japanese companies use big parallel machines now. The Japanese firms are more reluctant than their American counterparts for several reasons. First, the recession in Japan is causing companies to cut back on R&D expenditures and new computing systems are not spared in this regard. At the present time there is only one important Japanese manufacturer of parallel processing equipment; however, that company's product would probably be regarded as very expensive, even in the best of economic times.

Second, in Japan there is a more conservative approach to new technologies in general, and companies that have been producing simulation results on vector computers are not willing to devote the effort to rewrite computer codes for parallel systems. To be sure, this attitude also prevails in many U.S. companies. An important difference, though, between the U.S. and Japan, is the enormous experience in parallel processing that has accumulated at U.S. National Laboratories and universities. This has translated into a much larger interest in parallel processing at U.S. companies, and therefore a distinct advantage over Japanese companies. Most importantly, there is a critical difference in the way that U.S. and Japanese companies are able to access state-of-the-art computing technologies. For example, an American company wishing to "dabble" in parallel processing may do so by forming partnerships with the Laboratories or NSF National Supercomputer Centers. In so doing they are able to determine, in a reduced-risk environment, which of the new technologies can benefit their applications. In contrast, in Japan there are few, if any, national computing resources that are available to industry. Even if there

were, experience with the latest machines is still lacking in Japan. Also, most facilities in Japan that have the latest machines are dedicated to a single kind of scientific research, so it is difficult for those outside that field to collaborate.

A key phrase in the preceding paragraph was "reduced-risk," which simply means being able to gain experience with parallel machines without having to buy one. A state-of-the-art supercomputer represents a huge investment, in terms of initial cost as well as in continuing hardware and software support. (Thus, although we have not studied the correlation in detail, it seems clear that only the biggest companies, both in the U.S. and Japan, can afford to buy supercomputers.) Although high-performance computing is now in a state where enormous benefits can be reaped from use of the most recent hardware, these benefits generally do not come without significant expenditure of effort in tuning programs, understanding their performance, and choosing the right machine for a given application from a wide variety of choices. The situation is quite different from that of about 15 years ago, when vector processors started to replace mainframes and non-vector supercomputers. In that case, it was generally much easier to achieve an increase in performance with a new machine, and it was also easier, before using the machine, to estimate the improvement it would provide.

Another aspect of the risk involved with parallel computing relates back to applications software. Companies that develop applications must try to target their products to both hardware platforms that have large user bases and companies that have reasonable expectation for long-term survival. With so many different strategies for designing parallel systems, with so little standardization of parallel programming languages, and with so many relatively small, capital-poor companies designing parallel systems, the parallel computing market is in a state of flux and great uncertainty.

Thus, we believe that availability of large, multi-disciplinary computing research centers are an important advantage for both U.S. companies as users and U.S. supercomputer vendors, and that this attests to the success of DOE and NSF high-performance computing programs. In our survey we found several examples where companies in Japan will own their own supercomputers whereas companies of the same kind in the U.S. do not own their own machines. However, in several instances, the U.S. companies carry out the same kind of simulation using an NSF or DOE machine. We also found examples in Japan of companies that wanted to experiment with parallel processing but had no means of doing so.

Note however, that by reducing risk for American companies, the U.S. government remains as the primary funding source for advanced parallel systems. The high cost of most supercomputers makes their very existence a pawn of other factors, such as government support and acceptance. This also has played a large part in some recent upheavals in the supercomputer business during the past year.

In both the U.S. and Japan the real "grand computing challenge" is development of portable and reliable applications software for parallel machines. So far there is a dearth of applications of interest to commercial users. Thus, although there are currently many more companies that have installed large-scale parallel machines in the U.S. than there are in Japan, it is not yet clear that parallel processing has

provided a significant competitive advantage in the commercial sector. The few select exceptions to this generally do not involve numerical simulations; rather, they involve manipulating large quantities of data. The majority of parallel computers installed in companies are still primarily being tested and developed and have not contributed significantly to "getting there first." Vector processors remain the production computing workhorse.

The Changing Face of Supercomputing

Given the additional risk in successful use of today's newest supercomputers and the known large startup costs, the question that arises is who, if anyone, will continue to be able to afford them? As a discipline, "supercomputational" research, meaning the kind of research one can do using a supercomputer, is definitely here to stay, having proven both its economic as well as its intellectual value. However, it is likely that in the future, more and more supercomputer simulations will not be carried out using what we now call supercomputers.

As is well known, advances in microprocessor performance and semiconductor integration levels mean that much less expensive computing solutions such as workstations are attractive for scientific computing. We can see few reasons why the popularity of workstations (and personal computers) will not continue to grow. For example, a common practice in supercomputing today is to set up a simulation using high-level graphics and Computer Aided Design (CAD) systems running on a workstation. Then the simulation input is sent over a network to a supercomputer, where the actual "number crunching" takes place. Afterwards, the results are sent back, and the graphical output of the simulation is viewed on the workstation. As workstations become more and more powerful, there may be little reason to be encumbered with the supercomputer step at all.

Even today, although supercomputers still maintain significant advantages over workstations in raw speed and memory size, many users have already eliminated supercomputing from their R&D efforts. An important example is the oil industry, one of the largest and wealthiest of supercomputer users. Western Atlas Software, a company that provides software to oil industry users recently reported that new software is not being developed with vector computers in mind. Western Atlas finds that its customers are quite willing to accept the longer waiting time for their simulations to complete on less expensive hardware. Other factors, such as the desire to maintain separate, autonomous computing facilities at exploration sites, rather than a single, centralized one located at a remote research laboratory, are causing petroleum companies to abandon expensive large-scale computing hardware.

This is not to say that there will be absolutely no role for the biggest machines. Again, the petroleum industry is a good example, having recently purchased some of the biggest configurations of the newest massively parallel machines and using them in time-critical applications of corporate-wide importance. Also, the migration away from vector supercomputers may be counterbalanced by the discovery of new applications that can only be carried out on the most powerful supercomputers. The emergence of

datamining applications in the last year or so is an example. However, over time, workstations will continue to absorb more and more of existing supercomputer workloads. Vector and parallel supercomputers will not disappear, but their unique range of application will diminish substantially.

In fact, in the United States, we are willing to venture the guess that in the last year or so more progress in the application of high-performance computing to science and engineering problems has come from advances in areas unrelated *per se* to the supercomputer itself; rather, they have come from advances in moving, visualizing, storing, and sharing data.

Note that workstations will provide high-performance computing to a greater number of people, in effect, providing "supercomputing for the masses." It is possible that such a trend may have a greater effect in Japan than in the U.S., because up to now, supercomputing has been less available to large numbers of users in Japan.

We have some general comments about this report and the research that led to it. An "assessment of supercomputing usage" turned out to be a larger and more complicated task than we had first envisioned. The field is so dynamic that it is difficult to provide an assessment that accurately reflects the situation at both the beginning and the end of the study. This aspect is particularly true for the kinds of analyses given in Chapter 2. Also, the breadth of the supercomputing field today, in geographic as well as scientific reach, virtually ensures that a fully comprehensive study is difficult at best. Finally, although field research is an invaluable means of data collection, one must be careful not to over-generalize from the results and to note any inherent biases the researchers may have. In fact, this study showed how even three observers from the same home institution could easily reach different conclusions based on the same observations. There is, however, one clear result on which we all agree and it is the reason that all of us continue to work in this area: High-performance computing remains one of the most exciting and important developments of our time, and one that will continue to provide the opportunity for immense achievement in both science and commerce. As Richard Feynman said, "This isn't like driving down Route 66 and stopping at a Holiday Inn; this is an adventure!"

References

Chapter 1.

1. Kaufmann, W. J. and Smarr, L. L., "Supercomputing and the Transformation of Science," W. H. Freeman (New York), 1993.
2. Wasserman, H. J., Simmons, M. L., Bucher, I. Y., and Hiromoto, R. E., "An Assessment of High-Performance Computer Technology in Japan and Europe," Los Alamos National Laboratory Unclassified Release LA-UR-93-3604, December, 1993.
3. Sha, L., "Industrial Computing Applications," IEEE Computer, vol. 27, No. 1, 1994, page 13.
4. Furutsuki, T., "The Japanese Supercomputing Market, Part I. Opportunities in Japan for U.S. Vendors," Supercomputing Review, July, 1991, pp 34-36.

Chapter 2.

1. Dongarra, J. D., Meuer, H. W., and Strohmaier, E., "Top500 Supercomputer Report 1993," University of Mannheim Computing Center Publication RUM 33/93, <ftp://uni-mannheim.de/info/rumdoc/top500.ps.Z>, November, 1993.
2. Kahaner, D. K., "List of Supercomputers in Japan," file j-supers on University of Arizona Department of Computer Science anonymous ftp archive [cs.arizona.edu](ftp://cs.arizona.edu), 1990.
3. Kahaner, D. K., "Supercomputers in Japan," file j-supers.91 on University of Arizona Department of Computer Science anonymous ftp archive [cs.arizona.edu](ftp://cs.arizona.edu), 1992.
4. Kahaner, D. K., "Installed Super/Parallel Computers in Japan," file j-supers.92 on University of Arizona Department of Computer Science anonymous ftp archive [cs.arizona.edu](ftp://cs.arizona.edu), December, 1992.
5. Johnson, G. M. and Cavallini, J. S., "Grand Challenges, High-Performance Computing, and Computational Science," Singapore Supercomputing Conference '90 Proceedings, World Scientific Publishing (Singapore), 1991, pp 28-38.
6. Dongarra, J. J., "Performance of Various Computers Using Standard Linear Equations Software, available from anonymous ftp archive [ftp.ornl.gov](ftp://ftp.ornl.gov), February, 1994.
7. Lubeck, O. M., "Supercomputer Performance: The Theory, Practice, and Results," Los Alamos National Laboratory Report LA-11204-MS, 1988.
8. Wasserman, H. J., Simmons, M. L., Bucher, I. Y., and Hiromoto, R. E., "An Assessment of High-Performance Computer Technology in Japan and Europe," Los Alamos National Laboratory Unclassified Release LA-UR-93-3604, December, 1993.

9. Rene Haussamman, NEC Corporation, Manno, Switzerland, personal communication to hjw, February, 1994.
10. Kahaner, D. K., "Numerical Wind Tunnel at Japan's National Aerospace Lab," file nwt.793 on University of Arizona Department of Computer Science anonymous ftp archive cs.arizona.edu., July, 1993.
11. Kahaner, D. K., "Int Conf on Supercomputing 1993 & High Perf Comput Benchmk WS," file ics.93 on University of Arizona Department of Computer Science anonymous ftp archive cs.arizona.edu., August, 1993.
12. Furutsuki, T., "The Japanese Supercomputing Market, Part I. Opportunities in Japan for U.S. Vendors," *Supercomputing Review*, July, 1991, pp 34-36.
13. Ruediger Rohrbach, SIEMENS NIXDORF INFORMATIONSSYSTEME AG, Munich, Germany, personal communication to hjw, February, 1994.
14. Smaby, G., "Tracking Supercomputer Vendors: A Cakewalk, No More," *High Performance Computing Wire Electronic News Service Message #334*, May 6, 1993.
15. *High Performance Computing Wire Electronic News Service Message #121*, February 26, 1994.
16. Talwani, M., "3-D Seismic Imaging at the Houston Advanced Research Center: Using the SX-3 System in the Energy Industry," *NEC SX WORLD*, No. 14, Spring, 1994, pp 3-4.
17. Schutzer, D., "Parallel Processing and the Future Data Center: Computing in the Land of the Lilliputians, Van Nostrand Reinhold: New York, 1994.
18. Herbst, K., "Supers and Secrecy on Wall Street: Crunching Big Numbers for Big Bucks," *Supercomputing Review*, January, 1991, pp 26-29.
19. Kaji, T., "SX-1A at Daiwa Institute of Research Contributes to System Operation in the Financial and Capital Markets," *NEC SX WORLD*, No. 7, Spring, 1991, pp 5-6.
20. Jones, R. D., "Machines that Learn: Adaptive Computation with Neural Nets," *Los Alamos Science*, Los Alamos National Laboratory, Number 21, 1993, pp 195-203.
21. Faber, V., Hochberg, J. G., Kelly, P. M., Thomas, T. R., and White, J. M., "Concept Extraction a Data-Mining Technique," *Los Alamos Science*, Los Alamos National Laboratory, Number 22, 1994, pp 122-137.
22. Carlson, E., "MPP for the Commercial World: One User's Experience," *High Performance Computing Wire Electronic News Service Message #2852*, Nov. 5, 1993.
23. McCormack, R., "The Paragon Paradox," *High Performance Computing and Communications Week*, Vol. 3, Number 14, April 7, 1994.
24. Zanardelli, C., "The Emerging Need for Supercomputer-Trained Engineers at Ford," in *Supercomputers: A Key to U.S. Scientific, Technological, and Industrial Preeminence*, J. R. Kirkland, Ed., Praeger (New York), 1987.
25. Dodd, G. D., "Productivity, Cost, and Quality: Supercomputers at General Motors," in *Supercomputers: A Key to U.S. Scientific, Technological, and Industrial Preeminence*, J. R. Kirkland, Ed., Praeger (New York), 1987.

26. Morneau, H. P., "Increasing Business Value and Enhancing Industrial Competitiveness with Modeling and Simulation," Proc. Intl. Forum on Managing Computational Science for Industrial Competitiveness Conf., Swiss Scientific Computing Center, Feb., 1994.
27. Thomas, R. J., "Supercomputing at DuPont: A Chemical Engineering Perspective, CRAY CHANNELS, Cray Research Inc., Mendota Heights, MN, Vol. 13, No. 3, 1991, pp 18-21.
28. Dixon, D. A., Dobbs, K. D. and Walker, S., "Computational Chemical Design for CFC Alternatives at DuPont," CRAY CHANNELS, Cray Research Inc., Mendota Heights, MN, Vol. 16, No. 1, 1994, pp 7-11.
29. Pensak, D., "Training Supercomputer Scientists at DuPont," in Supercomputers: A Key to U.S. Scientific, Technological, and Industrial Preeminence, J. R. Kirkland, Ed., Praeger (New York), 1987.
30. Richards, W. G., "Supercomputers in Drug Design," in Supercomputational Science, R. G. Evans and S. Wilson, eds, Plenum Press (New York), 1990, pp 269-273.
31. Brull, L., Lang, L., Zeller, R., and Zitney, S., "Bayer AG and Cray Research Collaborate on Plantwide Dynamic Process Simulations," CRAY CHANNELS, Cray Research Inc., Mendota Heights, MN, Vol. 16, No. 1, 1994, pp 1-6.
32. Focardi, S., "Case Studies in Supercomputer Usage," Proc. Intl. Forum on Managing Computational Science for Industrial Competitiveness Conf., Swiss Scientific Computing Center, Feb., 1994.
33. Weber, H., and Widmer, A., "Experiences on Sandoz's Alliant FX-2800," Centro Svizzero di Calcolo Scientifico "Speedup", Volume 6, No 2, 1992, pp 67-68.
34. High Performance Computing Wire Electronic News Service Message #121, June 22, 1993.
35. Kahaner, D. K., "The Advanced Telecommunication Research Institute," file atr.92 on University of Arizona Department of Computer Science anonymous ftp archive cs.arizona.edu., February, 1992.
36. Kahaner, D. K., "Supercomputing at Institute for Materials Research," file imr-a.94 on University of Arizona Department of Computer Science anonymous ftp archive cs.arizona.edu., June, 1994.
37. Herbst, K., "Computational Fluid Dynamics in Japan," Supercomputing Review, January, 1992.
38. "Fast Breeder Reactor Program Moving Forward," in Science and Technology Perspectives, Foreign Broadcast Information Service, Vol. 9, No. 4, June, 1994, pp 10-15.
39. Samuels, R. J., "Japan's Aerospace Industry and Technology Ideology," Japan Digest, Spring, 1993, pp 21-24.
40. Saito, F., "Computational Science in Basic Research in an Electronics Industry," Proc. Intl. Forum on Managing Computational Science for Industrial Competitiveness Conf., Swiss Scientific Computing Center, Feb., 1994.
41. Kahaner, D., "Fujitsu Parallel Computing Workshop," file pcw.94 on University of Arizona Department of Computer Science anonymous ftp archive cs.arizona.edu., November, 1994.
42. Izumikawa, T., "Use of SX System at Mazda," NEC SX WORLD, No. 4, Summer, 1991, pp 5-6.
43. Kondoh, T., "The Use of the SX System at Toyota Central Research and Development Laboratories, Inc.," NEC SX WORLD, No. 11, Spring, 1993, pp 5-6.

44. Himeno, R. "Numerical Analysis and Visualization of Flow in Automobile Aerodynamics Development," Proc. Supercomputing Japan 91, Meridian Pacific Group (Mill Valley, CA), 1991, pp 52-57.
45. Himeno, R., Fujitani, K., and Takagi, M., "Simulation of Unsteady Three-Dimensional Flows Around an Automobile and its Visualization," Supercomputing in Fluid Flow, T.K.S. Murthy, C.A. Brebbia, editors, Computational Mechanics Publications (Southampton); Computational Mechanics Publications, 1993, pp 127-146.
46. Zaphyr, P., "The Westinghouse Network," in Supercomputers: A Key to U.S. Scientific, Technological, and Industrial Preeminence, J. R. Kirkland, Ed., Praeger (New York), 1987.
47. Omura, Y., "Some Results of CAE Applications at MHI Kobe," Proc. CUG Fall 1993, Cray User Group, Inc., Eagan, MN, pp 98-109.
48. Yanase, M. "Use of the SX-3/1LR System at Sumitomo Rubber Industries, Ltd." NEC SX WORLD, No. 13, Winter, 1994, pp 5-6.
49. Guillet, A., "Supercomputing and Tire Design at Michelin," CRAY CHANNELS, Cray Research Inc., Mendota Heights, MN, Vol. 10, No. 4, 1989, pp 10-13.
50. Yoshida, Y., Oka, A., and Iizuka, K., "Interactive Device Simulation at Toshiba," CRAY CHANNELS, Cray Research Inc., Mendota Heights, MN, Vol. 13, No. 2, 1991, pp 21-23.
51. Kahaner, D. K., "Japanese High Performance Computing Overview," file j-hpc.93 on University of Arizona Department of Computer Science anonymous ftp archive cs.arizona.edu., December, 1993.
52. Guinn, B., "Apple Computer PowerBook Computer Molding Simulation, CRAY CHANNELS, Cray Research Inc., Mendota Heights, MN, Vol. 14, No. 1, 1992, pp 10-12.
53. Toyofuku, T. "SX-3 System at Sanyo Electric," NEC SX WORLD, No. 8, Summer, 1992, pp 5-6.
54. Kahaner, D. K., "Computer Modeling in the Construction Industry--Ohbayashi Corporation's Research Laboratory," file ohbayashi on University of Arizona Department of Computer Science anonymous ftp archive cs.arizona.edu, May, 1990.
55. Omote, Y., "Use of the SX-3 System at the Ohbayashi Corporation," NEC SX WORLD, No. 14, Spring, 1994, pp 5-6.
56. Japan Company Handbook, Toyo Keizai (Bloomington, MN), Spring 1994 Edition.
57. Maisel, M., "Industrial Engineering-University Collaboration: A New Mold," CRAY CHANNELS, Cray Research Inc., Mendota Heights, MN, Vol. 12, No. 2, 1990, pp 9-11.
58. Dilber, I. and Cheng, W., "The Impact of Synergy: An Example From the NCSA-FMC Partnership," CRAY CHANNELS, Cray Research Inc., Mendota Heights, MN, Vol. 13, No. 2, 1991, pp 12-13.
59. Farrar, C.R., et al., "Dynamic Characterization and Damage Detection in the I-40 Bridge over the Rio Grande," Los Alamos National Laboratory Manuscript LA-12767-MS, June, 1994.
60. Schmitt, G., "The Role of Communication in Design and Construction Industry Management," Proc. Intl. Forum on Managing Computational Science for Industrial Competitiveness Conf., Swiss Scientific Computing Center, Feb., 1994.

61. Kahaner, D. K., "Computer Related Activities at NKK," file nkk.93 on University of Arizona Department of Computer Science anonymous ftp archive cs.arizona.edu, September, 1993.
62. Kahaner, D. K., "Real World Computing Summary by Nikkei Computer," file rwk.693 on University of Arizona Department of Computer Science anonymous ftp archive cs.arizona.edu, July, 1993.
63. Makino, T., "SX System in Sumitomo Metal Industries," NEC SX WORLD, No. 6, Winter, 1992, pp 5-6.
64. Schneider, M., "Modeling Steel Metallurgy in Continuous Casting Tundishes," High Performance Computing Wire Electronic News Service Message #1392, April 16, 1994.
65. Demuyne, M., and Bedu, P., "Network Supercomputing with Database Management at EDF," CRAY CHANNELS, Cray Research Inc., Mendota Heights, MN, Vol. 14, No. 1, 1992, pp 13-15.
66. Kahaner, D. K., "Supercomputer Sales to Japanese Government Facilities," file sc-new.93 on University of Arizona Department of Computer Science anonymous ftp archive cs.arizona.edu, May 1993
67. Kahaner, D. K., "Re: Japanese Supercomputer Purchases," file j-sc-b.94 on University of Arizona Department of Computer Science anonymous ftp archive cs.arizona.edu, May, 1994.
68. Kahaner, D. K., personal communication, 13 April 1993.
69. "Who Says Science Has to Pay Off Fast?" Business Week, March 21, 1994, pp 110-111.
70. Comments by Yasushi Takeda, Hitachi, Ltd. Managing Director, in "Electronics Manufacturers Review R&D Projects," Foreign Broadcast Information Service JPRS-JST-94-010, March, 1994.
71. "STA to Start Computational Science Project in FY95," Foreign Broadcast Information Service JPRS-JST-94-030, September, 1994.
72. "U.S. Monitoring Japan's Supercomputer Buys Before Action," High Performance Computing Wire Electronic News Service Message #1019, May 4, 1993.
73. Fatoohi, Raad, "Visit to Real World Computing Research Center," file ics.93 on University of Arizona Department of Computer Science ftp archive cs.arizona.edu, August, 1993.
74. Anchordoguy, M., "Japan's Strategic Trade and Industrial Policies, A Thing of the Past?" Political Science Quarterly, to appear, 1994.

Chapter 3.

1. Focardi, S., "Case Studies in Supercomputer Usage," Proc. Intl. Forum on Managing Computational Science for Industrial Competitiveness Conf., Swiss Scientific Computing Center, Feb., 1994.
2. Japan Company Handbook, Toyo Keizai (Bloomington, MN), Spring 1994 Edition.
3. Kahaner, D. K., "International Joint Conference on Neural Nets," file ijcn.93 on University of Arizona Department of Computer Science anonymous ftp archive cs.arizona.edu, December, 1993.

Chapter 4.

1. Wasserman, H. J., Simmons, M. L., Bucher, I. Y., and Hiromoto, R. E., "An Assessment of High-Performance Computer Technology in Japan and Europe," Los Alamos National Laboratory Unclassified Release LA-UR-93-3604, December, 1993.
2. National Research Council, "Computing the Future A Broader Agenda for Computer Science and Engineering," National Academy Press, 1992.
3. Office of Science and Technology Policy, "High Performance Computing and Communications Toward a National Infrastructure," Report by the Committee on Physical, Mathematical, and Engineering Sciences, 1994.

Glossary

Algorithm:	A specific set of steps used to solve a computational problem.
Architecture:	General term that describes the design features of a computer system.
Benchmark:	The process by which the true computational speed of a computer is measured.
CFD	Computational Fluid Dynamics; using numerical methods to simulate the flow of gases and liquids.
Chip:	An electronic device made of semiconductor material, usually silicon, on which digital logic circuits or digital memory circuits are printed.
Clock cycle:	All computers have their own internal clock. The clock cycle is the period of time between clock ticks. Also known as clock period, it is quoted generally in nanoseconds (10^{-9} seconds), and is the reciprocal of clock frequency (or clock rate) which is generally quoted in megahertz (MHz).
Compiler:	Computer software that translates a user's program into machine language.
Connection Machine:	A series of massively parallel computer systems manufactured by Thinking Machines Inc. (Cambridge, MA).
GFLOPS:	"Giga-flops," or billions of floating-point operations per second, a unit of computer performance for scientific calculations.
IC:	integrated circuit.
LINPACK	A mathematical analysis software package. One of the routines included is frequently timed and used as a "benchmark" performance metric.
MHz:	Megahertz, or millions of cycles per second.
MFLOPS:	"Mega-flops," or millions of floating-point operations per second, a unit of computer performance for scientific calculations.
Microprocessor:	A central processing unit implemented on a small set of (generally 1-8) computer chips.
MIMD:	Multiple-instruction multiple-data, a programming model in which all processors in a system execute their own instruction streams asynchronously.
MPP:	Massively parallel processor, a type of computer architecture employing at least hundreds, and generally thousands of processors.
Multiprocessor:	Any computer system with more than one processor.

Network:	In a multiprocessor system, the hardware used to interconnect processors and allow communication between them. Also referred to as interconnection network or communication network. Networks are characterized by, among other things, their topology, latency, and bandwidth.
Node:	Synonym for processor or processing element.
NWT	A parallel processing supercomputer designed by Fujitsu installed at Japan's National Aerospace Laboratory
Paragon	A series of massively-parallel supercomputers manufactured by Intel.
Pipeline:	Term used to refer to the series of consecutive hardware steps that must be carried out to perform arithmetic in most computers.
RISC:	Reduced instruction set computer. A term generally used to describe the kinds of microprocessors used in scientific workstations.
Scalar:	Refers to a single data value, as opposed to a string of similar values. Often used to refer to computation that cannot be vectorized.
Scientific workstation:	A relatively small, inexpensive computer system generally consisting of a fast microprocessor, disk, and a large video display terminal.
Shared memory:	A type of memory architecture in which any location in the memory is accessible from any processor in the system.
SIMD:	Single-instruction multiple-data, a programming model in which all processors in a system have to execute the same instruction on data located in their memories at identical locations.
SPARC:	Scalable processor architecture, a proprietary microprocessor design owned by Sun Microsystems, Inc.
Superscalar:	A computer architecture in which more than one instruction is issued in a given clock cycle.
T3D	A series of massively-parallel supercomputers manufactured by Cray Research, Inc.
TFLOPS:	"Tera-flops," or trillions of floating-point operations per second, a unit of computer performance for scientific calculations.
Vector computer:	A computer comprised of one or more vector processors.
Vector processor:	A processor that carries out computation in an assembly line fashion, and generates more than one result with a single instruction.
Vectorize:	To cause a computation to be performed in using vector hardware.
VPP500	A parallel supercomputer manufactured by Fujitsu, Ltd.
Workstation:	Same as scientific workstation.

1961-1962

1961-1962

LOS ALAMOS NAT'L LAB.
LIB. REPT. COLLECTION
RECEIVED

'95 MAY 22 AM 3 31